HyperFabric Administrator's Guide

HP-UX 11i v2

Edition 14

Manufacturing Part Number : B6257-90043
March 2004

United States

Legal Notices

The information in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Warranty

A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

U.S. Government License

Proprietary computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

Trademark Notices

Oracle ® is a registered trademark of Oracle Corporation.

UNIX® is a registered trademark in the United States and other countries, licensed exclusively through The Open Group.

Contents

Contents

5. Managing HyperFabric

6. Troubleshooting HyperFabric

Contents

Figures

Figures

Tables

Tables

About This Document

This document describes how to install, configure, and troubleshoot the HyperFabric product on the HP-UX 11i v2 (HP-UX 11.23) operating system.

Before you install the HyperFabric software, ensure that the operating system software and the appropriate files, scripts, subsets are installed.

The document printing date and part number indicate the document's current edition. The printing date will change when a new edition is printed. Minor changes may be made at reprint without changing the printing date. The document part number will change when extensive changes are made.

Document updates may be issued between editions to correct errors or document product changes. To ensure that you receive the updated or new editions, you should subscribe to the appropriate product support service. See your HP sales representative for details.

The latest version of this manual can be found online at http://www.docs.hp.com/hpux/netcom/index.html#HyperFabric.

Intended Audience

This manual is intended for system and network administrators responsible for installing, configuring, and managing the HyperFabric software and hardware. Administrators are expected to have knowledge of operating system concepts, commands, and configuration.

It is helpful to have knowledge of Transmission Control Protocol/Internet Protocol (TCP/IP) networking concepts and network configuration.

This document is not a tutorial.

HP-UX Release Name and Release Identifier

Each HP-UX 11i release has an associated release name and release identifier. The *uname* (1) command with the `-r` option returns the release identifier. Table 1 shows the releases available for HP-UX 11i.

Table 1 HP-UX 11i Releases

Release Identifier	Release Name	Supported Processor Architecture
B.11.11	HP-UX 11i v1	PA-RISC
B.11.20	HP-UX 11i v1.5	Intel® Itanium®
B.11.22	HP-UX 11i v1.6	Intel® Itanium®
B.11.23	HP-UX 11i v2	Intel® Itanium®

Publishing History

Table 2 provides the publication date for the pertinent edition number.

Table 2 Publishing History Details

Edition Number	Publication Date
First	March 1998
Second	June 1998
Third	August 1998
Fourth	October 1998
Fifth	December 1998
Sixth	February 1999
Seventh	April 1999
Eighth	March 2000
Ninth	June 2000
Tenth	December 2000

Table 2 Publishing History Details (Continued)

Edition Number	Publication Date
Eleventh	June 2001
Twelfth	September 2002
Thirteenth	July 2003
Fourteenth	March 2004

What Is in This Document

HyperFabric Administrator's Guide is divided into several chapters, each of which contains information about installing, configuring, or troubleshooting HyperFabric. The appendixes contain supplemental information.

The following list describes the content in more detail.

Table 3 Organization

Chapter	Description
Overview of the HyperFabric Product	Presents an overview of HyperFabric and lists the components that the HyperFabric product contains
Planning the Fabric	Describes the steps that need to be followed while planning the fabric
Installing HyperFabric	Describes the tasks to install the HyperFabric products on the HP-UX 11i v2 operating system
Configuring HyperFabric	Describes the tasks to configure HyperFabric
Managing HyperFabric	Describes the tasks to start, stop and manage HyperFabric
Troubleshooting HyperFabric	Describes how to troubleshoot HyperFabric

New and Changed Information in This Edition

This edition includes information about the transparent local failover feature of Hyper Messaging Protocol (HMP). This feature is available with the HyperFabric version B.11.23.01.

Typographic Conventions

This document uses the following typographic conventions:

Book Title	Italic (slanted) type indicates document and book names.
`daemon`	Courier font type indicates daemons, files, commands, manpages, and option names.
[ ] { }	In syntax definitions, square brackets indicate items that are optional and braces indicate items that are required.

HP Welcomes Your Feedback

HP welcomes any comments and suggestions you have on this manual.

You can send your comments in the following ways:

- Internet electronic mail: `netinfo_feedback@cup.hp.com`
- Using a feedback form located at the following URL:

  `http://docs.hp.com/assistance/feedback.html`

Please include the following information along with your comments:

- The complete title of the manual and the part number. (The part number appears on the title page of printed and PDF versions of a manual.)
- The section numbers and page numbers of the information on which you are commenting.
- The version of HP-UX that you are using.

Please note that the HP-UX networking communications publications group does not provide technical support for HP products.

1 Overview of the HyperFabric Product

This chapter contains the following sections that give general information about HyperFabric:

About HyperFabric

HyperFabric is an HP high-speed, packet-based interconnect for node-to-node communications. HyperFabric provides higher speed, lower network latency and uses less CPU than other industry standard protocols (for example, Fibre Channel and Gigabit Ethernet). Instead of using a traditional bus-based technology, HyperFabric is built around switched fabric architecture, providing the bandwidth necessary for high speed data transfer. This clustering solution delivers the performance, scalability, and high availability required by the following:

- Parallel Database Clusters
  - — Oracle 9i Real Application Clusters (RAC)
  - — Oracle 8i Parallel Servers (OPS)
- Parallel Computing Clusters
- Client/Server Architecture Interconnects (for example, SAP)
- Multi-Server Batch Applications (for example, SAS Systems)
- Enterprise Resource Planning (ERP)
- Technical Computing Clusters
- HP Message Passing Interface (MPI) based applications
- OpenView Data Protector (earlier known as Omniback)
- Network Backup
- Data Center Network Consolidation
- E-services

HyperFabric Products

HyperFabric hardware consists of host-based interface adapter cards, interconnect cables, and optional switches. HyperFabric software resides in Application Specific Integrated Circuits (ASICs) and firmware on the adapter cards and includes user-space components and HP-UX drivers.

Currently, fiber-based HyperFabric hardware are available. In addition, a hybrid switch that has 8-fiber ports is available to support HF2 clusters.

This section describes the various HyperFabric products. For more information on HP 9000 systems that support HyperFabric products, see the *HyperFabric Release Notes,* available at http://docs.hp.com/hpux/netcom/index.html#HyperFabric.

NOTE

This document uses the term **HyperFabric** (**HF**) to refer to the hardware and software that form the HyperFabric cluster interconnect product.

The term **HyperFabric2** (**HF2**) refers to the following fiber-based hardware components:

- The A6386A adapter
- The A6384A switch chassis
- The A6388A and A6389A switch modules. (Although the A6389A switch module has 4-copper ports, it is still considered an HF2 component because it can only be used with the A6384A HF2 switch chassis).
- The C7524A, C7525A, C7526A, and C7527A cables

HyperFabric Adapters

The HyperFabric adapters are as follows:

- A6386A HF2 PCI (4X) adapter with a fiber interface.

The A6092A HyperFabric adapter is supported beginning with the following HyperFabric software versions:

- HP-UX 11.0: HyperFabric software version B.11.00.09
- HP-UX 11i v1: HyperFabric software version B.11.11.00
- HP-UX 11i v2: HyperFabric software version B.11.23.00

The A6386A HyperFabric2 adapter is supported beginning with the following HyperFabric software versions:

- HP-UX 11.0: HyperFabric software version B.11.00.11
- HP-UX 11i v1: HyperFabric software version B.11.11.01
- HP-UX 11i v2: HyperFabric software version B.11.23.00

Switches and Switch Modules

The HyperFabric2 switches are as follows:

- A6384A HF2 fiber switch chassis with one integrated Ethernet management LAN adapter card, one integrated 8-port fiber card, and one expansion slot. For the chassis to be a functional switch, install one of the following switch modules in the expansion slot:
  - The A6388A HF2 8-port fiber switch module. This gives 16-fiber ports to the switch (8 from the integrated fiber card and 8 from the A6388A switch module).
  - The A6389A HF2 4-port copper switch module. This gives 12 ports to the switch - a mixture of 8-fiber ports (from the integrated fiber card) and 4-copper ports (from the A6389A module).

The A6384A HF2 switch chassis with either module installed is supported beginning with the following HyperFabric software versions:

- HP-UX 11.0: HyperFabric software version B.11.00.11
- HP-UX 11i v1: HyperFabric software version B.11.11.01
- HP-UX 11i v2: HyperFabric software version B.11.23.00

NOTE In this manual, the terms **HyperFabric2 switch** or **HF2 switch** refer to the functional switch (the A6384A switch chassis with one of the switch modules installed).

IMPORTANT HF2 adapters and switches are not supported by software versions earlier than those listed in “HyperFabric Adapters” on page 4 and “Switches and Switch Modules” on page 5.

To determine the version of HyperFabric, issue the following command:

```
$ swlist | grep -i hyperfabric
```

Other Product Elements

The following are the other elements of the HyperFabric product family:

- HF2 fiber cables
  - C7524A (2m length)
  - C7525A (16m length)
  - C7526A (50m length)
  - C7527A (200m length)
- The HyperFabric software: The software resides in ASICs and firmware on the adapter cards and includes user-space components and HP-UX drivers.

  HyperFabric supports the IP network protocol stack, specifically TCP/IP and UDP/IP.

  HyperFabric software includes HyperMessaging Protocol (HMP). HMP provides higher bandwidth, lower CPU overhead, and lower latency (the time a message takes to get from one point to another). However, these HMP benefits are available only when applications that are developed on top of HMP are running. HMP can only be used on HP 9000 systems running HP-UX 11.0 or 11i v1, provided HyperFabric A6092A or A6386A (PCI 4X) adapter cards are installed on those systems.

HyperFabric Concepts

This section briefly describes some of the basic HyperFabric concepts and terms.

The **fabric** is the physical configuration that consists of all the HyperFabric adapters, the HyperFabric switches (if any), and the HyperFabric cables connecting them. The network software controls data transfer over the fabric.

The HyperFabric configuration contains two or more HP 9000 systems and optional HyperFabric switches. Each HP 9000 acts as a **node** in the configuration. Each node has a minimum of one and a maximum of eight HyperFabric **adapters** installed in it. (For information on the maximum number of adapters that can be installed in each system, see Chapter 2, “Planning the Fabric,” on page 9.) Each HF2 switch can be configured with 12 or 16 ports. HyperFabric supports a maximum of 4 HyperFabric switches. You can mesh HyperFabric switches and configurations with up to four levels of meshed switches are supported.

You can plan a HyperFabric cluster as a **High Availability** (**HA**) configuration, when it is necessary to ensure that each node can always participate in the fabric. This is done by using ServiceGuard (earlier known as MC/ServiceGuard), ServiceGuard OPS Edition (earlier known as MC/LockManager), and the Event Monitoring Service (EMS). Configurations of up to 8 nodes are supported under ServiceGuard.

Beginning with HyperFabric software versions B.11.00.05 and B.11.11.00,you can use **relocatable IP addresses** as part of an HA configuration. Relocatable IP addresses permit a client application to reroute through an adapter on a remote node, allowing that application to continue processing without interruption. The rerouting is transparent. This function is associated with ServiceGuard (see “Configuring ServiceGuard for HyperFabric Relocatable IP Addresses” on page 95). When the monitor for HyperFabric detects a failure and the backup adapter takes over, the relocatable IP address is transparently migrated to the backup adapter. Throughout this migration process, the client application continues to execute normally.

When you start HyperFabric (with the `clic_start` command, through SAM, or by booting the HP 9000 system), you start the **management process**. This process must be active for HyperFabric to run. If the HyperFabric management process on a node stops running for some

reason (for example, if it is `killed`), all HyperFabric-related communications on that node are stopped immediately. This makes the node unreachable by other components in the fabric.

When you start HyperFabric, the fabric is verified automatically. This is because each node performs a self diagnosis and verification over each adapter installed in the node. In addition, the management process performs automatic routing and configuring for each switch (if switches are part of the fabric). You can, if needed, run the `clic_stat` command to get a textual map of the fabric, which can be used as another method of quick verification.

Notice that the commands to administer HyperFabric have a prefix of `clic_`, and some of the other components have **CLIC** as part of their name (for example, the CLIC firmware and the CLIC software). CLIC stands for CLuster InterConnect, and it is used to differentiate those HyperFabric commands or components from other commands or components. For example, the HyperFabric command `clic_init` is different from the HP-UX `init` command.

2 Planning the Fabric

This chapter contains the following sections that include general guidelines and protocol-specific considerations for planning HyperFabric clusters that run TCP/UDP/IP or HMP applications.

Preliminary Considerations

Before assembling a fabric physically, do the following to address all of the appropriate issues:

Step 1. Read Chapter 1, “Overview of the HyperFabric Product,” on page 1, to get a basic understanding of HyperFabric and its components.

Step 2. Read this chapter, *Planning the Fabric*, to gain an understanding of protocol specific configuration guidelines for TCP/UDP/IP and HMP applications.

Step 3. Read “Configuration Overview” on page 69, “Information You Need” on page 71, and “Configuration Information Example” on page 74, to understand how to configure the fabric.

Step 4. Decide the number of nodes that will be interconnected in the fabric.

Step 5. Decide the type of HP 9000 system for each node (for a list of supported HP 9000 systems, see the *HyperFabric Release Notes* available at

http://docs.hp.com/hpux/netcom/index.html#HyperFabric.

Step 6. Determine the network bandwidth requirements for each node.

Step 7. Determine the number of adapters needed for each node.

Step 8. Determine if a High Availability (ServiceGuard) configuration will be needed. If ServiceGuard is used, each node should have at least two adapters.

Step 9. Decide the topology of the fabric.

Step 10. Determine how many switches will be used based on the number of nodes in the fabric. The only configuration that can be supported without a switch is the node-to-node configuration (HA or non-HA). HyperFabric supports meshed switches up to a depth of four switches, starting with the following versions of the HyperFabric software:

- For HF2 switches: software versions B.11.00.11, B.11.11.01, and B.11.23.00.

Step 11. Draw the cable connections from each node to switches (if the fabric will contain switches). If you use an HA configuration with switches, it requires more than one switch for complete redundancy and to avoid a single point of failure. For example, each adapter can be connected to its own switch, or two switches can be connected to four adapters.

HyperFabric Features, Parameters and Supported Configurations for TCP/UDP/IP and HMP Applications

The following sections in this chapter define HyperFabric functionality for TCP/UDP/IP applications and Hyper Messaging Protocol (HMP) applications. There are distinct differences in supported hardware, available features and performance, depending on which protocol is used by applications running on the HyperFabric.

TCP/UDP/IP

TCP/UDP/IP is supported on all HF2 hardware. Although some of the HyperFabric adapter cards support both HMP and TCP/UDP/IP applications, in this section, the focus is on TCP/UDP/IP HyperFabric applications.

Application Availability

All applications, including Oracle 9i and HP-MPI, that use the TCP/UDP/IP stack are supported.

Features

This section discusses the following HyperFabric features on TCP/UDP/IP:

- **OnLine Addition and Replacement (OLAR): Supported**

  The OLAR feature allows the replacement or addition of HyperFabric adapter cards while the system (node) is running. HyperFabric supports this functionality on the SD64A, rx8620, rx4640, rp54xx (L-class), rp74xx (N-class), rp8400 and Superdome systems, running on the HP-UX 11i v2 platform.

  For more information on OLAR, including instructions for implementing this feature, see “Online Addition and Replacement” on page 44 and *Configuring HP-UX for Peripherals* Part Number B2355-90698 November 2000 Edition.

- **Event Monitoring Service (EMS): Supported**

  In the HyperFabric version B.11.23.01, the HyperFabric EMS monitor enables the system administrator to separately monitor each HyperFabric adapter on every node in the fabric, in addition to monitoring the entire HyperFabric subsystem. The monitor can inform the user if the resource being monitored is UP or DOWN. The administrator defines the condition to trigger a notification (usually a change in interface status). Notification can be accomplished with one of the following:

  — A Simple Network Management Protocol (SNMP) trap

— Logging into a user specified log file with a choice of severity

— Email to a user defined email address.

For more information on EMS, including instructions for implementing this feature, see “Configuring the HyperFabric EMS Monitor” on page 85 and the *EMS Hardware Monitors User’s Guide* Part Number B6191-90028 September 2001 Edition.

- **ServiceGuard: Supported**

  Within a cluster, ServiceGuard groups application services (individual HP-UX processes) into packages. In the event of a single service failure (node, network, or other resource), EMS provides notification and ServiceGuard transfers control of the package to another node in the cluster, allowing services to remain available with minimal interruption.

  ServiceGuard via EMS, directly monitors cluster nodes, LAN interfaces, and services (the individual processes within an application). ServiceGuard uses a heartbeat LAN to monitor the nodes in a cluster. ServiceGuard cannot use the HyperFabric interconnect as a heartbeat LAN. Instead, use a separate LAN for the heartbeat.

  For more information on configuring ServiceGuard, see “Configuring HyperFabric with ServiceGuard” on page 87, and *Managing MC/ServiceGuard* Part Number B3936-90065 March 2002 Edition.

- **High Availability (HA): Supported**

  To create a highly available HyperFabric cluster, there should not be any single point of failure. Once the HP 9000 nodes and the HyperFabric hardware have been configured with no single point of failure, ServiceGuard and EMS can be configured to monitor and fail over nodes and services using ServiceGuard packages.

  If any HyperFabric resource in a cluster fails (adapter card, cable or switch port), the HyperFabric driver transparently routes traffic over other available HyperFabric resources with no disruption of service.

  The ability of the HyperFabric driver to transparently fail over traffic reduces the complexity of configuring highly available clusters with ServiceGuard, because ServiceGuard has to take care of node and service failover only.

ServiceGuard uses a “heartbeat” to monitor the cluster. The HyperFabric links cannot be used for the heartbeat. Instead, an alternate LAN connection such as 100BaseT, Ethernet, Token Ring, or FDDI must be made between the nodes for use as a heartbeat link.

End-To-End HA: HyperFabric provides end-to-end HA on the entire cluster fabric at the link level. If any of the available routes in the fabric fails, HyperFabric transparently redirects all the traffic to a functional route and, if configured, notifies ServiceGuard or other enterprise management tools.

Active-Active HA: In configurations where there are multiple routes between nodes, the HyperFabric software uses a hashing function to determine an adapter or a route through which it sends messages. This is done on a message-by-message basis. All of the available HyperFabric resources in the fabric are used for communication.

In contrast to Active-Passive HA, where one set of resources is not utilized until another set fails, Active-Active HA provides the best return on investment because all of the resources are utilized simultaneously. ServiceGuard is not required for Active-Active HA operation.

For more information on setting up HA HyperFabric clusters, see Figure 2-3 “TCP/UDP/IP High Availability Switched Configuration” on page 24.

- **Dynamic Resource Utilization (DRU): Supported**

  If you add a new resource (node, adapter, cable or switch) to a cluster, the HyperFabric subsystem dynamically identifies the added resource and starts using it. The same process takes place when a resource is removed from a cluster. The difference between DRU and OLAR is that OLAR applies only to the addition or replacement of adapter cards from nodes.

- **Load Balancing: Supported**

  When an HP 9000 HyperFabric cluster is running TCP/UDP/IP applications, the HyperFabric driver balances the load across all available resources in the cluster, including nodes, adapter cards, links, and multiple links between switches.

- **Switch Management: Not Supported**

Switch Management is not supported. Switch management will not operate properly if you enable it on a HyperFabric cluster.

- **Diagnostics: Supported**

  Diagnostics can be run to obtain information on many of the HyperFabric components using the `clic_diag, clic_probe` and `clic_stat` commands, as well as the Support Tools Manager (STM).

  For more information on HyperFabric diagnostics, see "Running Diagnostics" on page 133.

Configuration Parameters

This section describes the maximum limits for TCP/UDP/IP HyperFabric configurations. There are numerous variables that can impact the performance of any HyperFabric configuration. For guidance on specific HyperFabric configurations for TCP/UDP/IP applications, see the section, "TCP/UDP/IP Supported Configurations" on page 21.

- HyperFabric is supported only on the HP 9000 series servers and workstations.
- TCP/UDP/IP is supported for all HyperFabric hardware and software.
- Maximum Supported Nodes and Adapter Cards

  In point-to-point configurations, the complexity and performance limitations of having a large number of nodes in a cluster make it necessary to include switching in the fabric. Typically, point-to-point configurations consist of only 2 or 3 nodes.

  In switched configurations, HyperFabric supports a maximum of 64 interconnected adapter cards.

  A maximum of 8 HyperFabric adapter cards are supported per instance of the HP-UX operating system. The actual number of adapter cards a particular node is able to accommodate also depends on slot availability and system resources. See node specific documentation for details.

  HyperFabric subsystem supports a maximum of 8 configured IP addresses per instance of the HP-UX operating system.

- Maximum Number of Switches

You can interconnect (mesh) up to 4 switches (16-port fiber, or Mixed 8 fiber ports) in a single HyperFabric cluster.

- Trunking Between Switches (multiple connections)

  You can use trunking between switches to increase bandwidth and cluster throughput. Trunking is also a way to eliminate a possible single point of failure. The number of trunked cables between nodes is limited only by port availability. To assess the effects of trunking on the performance of any particular HyperFabric configuration, contact your HP representative.

- Maximum Cable Lengths

  HF2 (fiber): The maximum distance is 200m (Four standard cable lengths are sold and supported: 2m, 16m, 50m and 200m).

  TCP/UDP/IP supports up to four HF2 switches connected in series with a maximum cable length of 200m between the switches and 200m between switches and nodes.

  TCP/UDP/IP supports up to four hybrid HF2 switches connected in series with a maximum cable length of 200m between fiber ports.

- Throughput and Latency

Table 2-1 HF2 Throughput and Latency with TCP/UDP/IP Applications

Server Class	Maximum Throughput	Latency
rp7400	2 + 2 Gbps full duplex per link	< 42 microsec

Table 2-2 Supported HyperFabric Adapter Configurations

HF Adapter	Bus Type	Supported HP Systems	HP-UX Version	OLAR Support?	Maximum Adapters per System
A6386A	PCI (4X)	rx2600 servers	11i v2	No	1
A6386A	PCI (4X)	rx56XX servers	11i v2	No	4
A6386A	PCI (4X)	zx6000 workstations	11i v2	No	1
A6386A	PCI (4X)	SD64A servers	11i v2	Yes	8 (maximum 4 per PCI card cage)
A6386A	PCI (4X)	rx7620 servers	11i v2	No	8 (maximum 4 per PCI card cage)
A6386A	PCI (4X)	rx8620 servers	11i v2	Yes	8 (maximum 4 per PCI card cage)
A6386A	PCI (4X)	rx4640 servers	11i v2	Yes	6

TCP/UDP/IP Supported Configurations

Multiple TCP/UDP/IP HyperFabric configurations are supported to match the cost, scaling, and performance requirements of each installation.

In the previous section, “Configuration Parameters” on page 17, the maximum limits for TCP/UDP/IP enabled HyperFabric hardware configurations were outlined. In this section the TCP/UDP/IP enabled HyperFabric configurations that HP supports are explained. These recommended configurations offer an optimal mix of performance and availability for a variety of operating environments.

There are many variables that can impact HyperFabric performance. If you are considering a configuration that is beyond the scope of the following HP supported configurations, contact your HP representative.

Point-to-Point Configurations

You can interconnect large servers like the HP Superdome to run Oracle RAC 9i and enterprise resource planning applications. These applications are typically consolidated on large servers.

Point-to-point connections between servers support the performance benefits of HMP without investing in HyperFabric switches. This is a good solution in small configurations where the benefits of a switched HyperFabric cluster might not be required (see configuration A and configuration C in Figure 2-1).

If there are multiple point-to-point connections between two nodes, traffic load is balanced over those links. If one link fails, the load fails over to the remaining links (see configuration B in Figure 2-1).

Running applications using TCP/UDP/IP on a HyperFabric cluster provides major performance benefits compared to other technologies such as Ethernet. If a HyperFabric cluster is originally set up to run enterprise applications using TCP/UDP/IP and the computing environment stabilizes with a requirement for higher performance, migration to HMP is always an option.

Figure 2-1 TCP/UDP/IP Point-To-Point Configurations

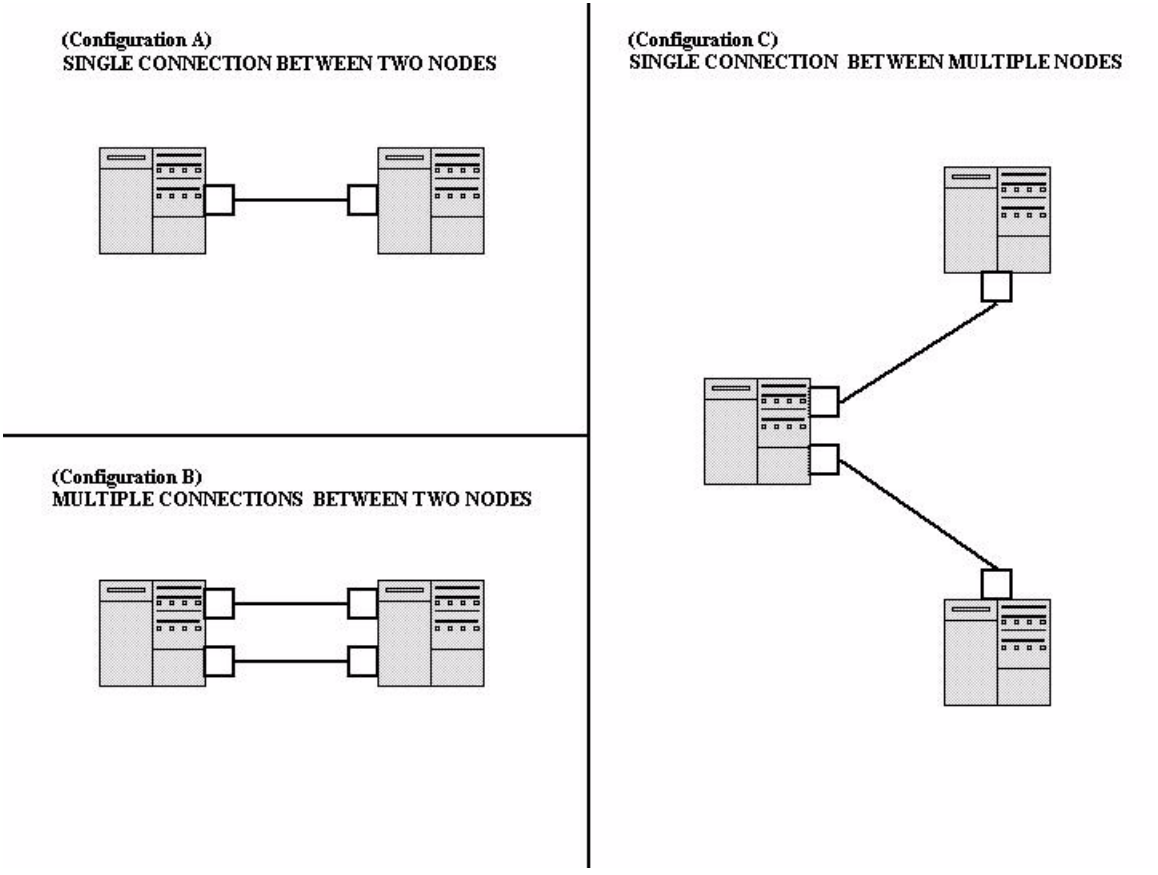

Switched Configuration

This configuration offers the same benefits as the point-to-point configurations illustrated in Figure 2-1, but it has the added advantage of greater connectivity (see Figure 2-2).

Figure 2-2 TCP/UDP/IP Basic Switched Configuration

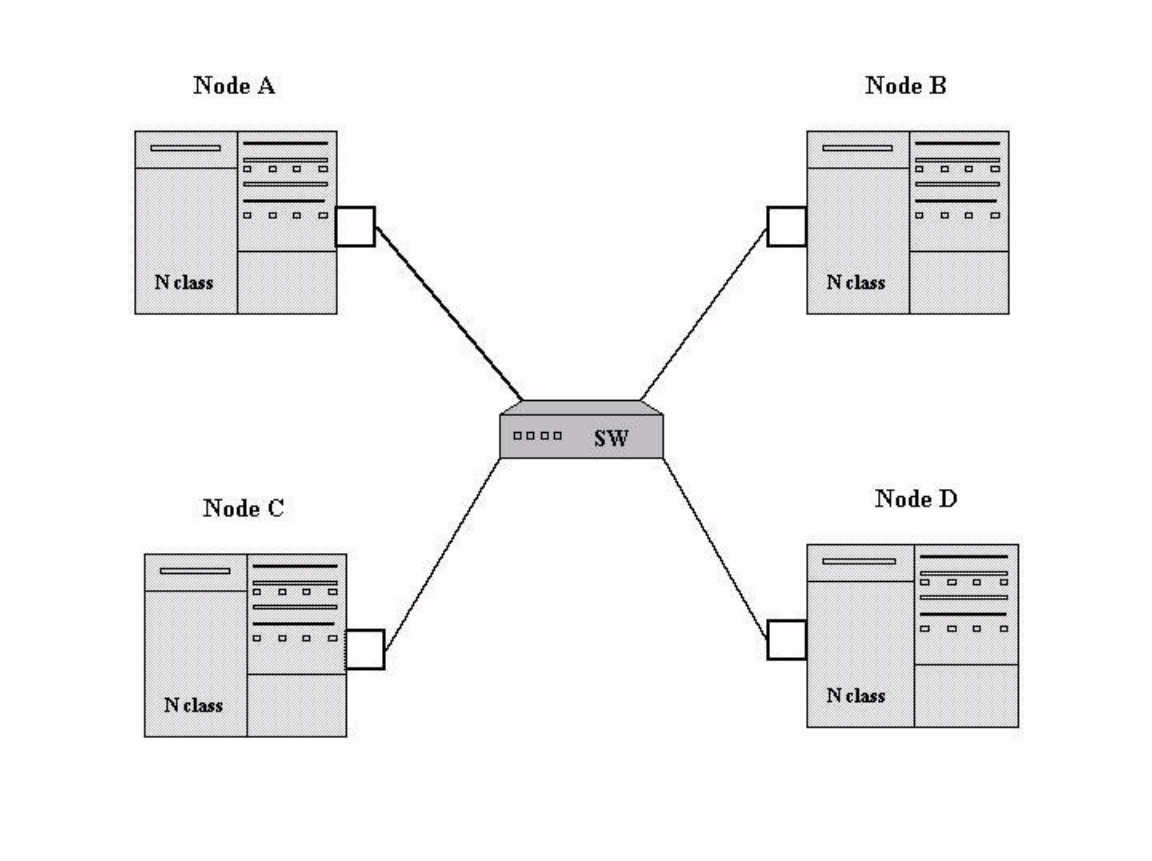

High Availability Switched Configuration

This configuration has no single point of failure. The HyperFabric driver provides end-to-end HA. If any HyperFabric resource in the cluster fails, traffic is transparently rerouted through other available resources. This configuration provides high performance and high availability (see Figure 2-3).

Figure 2-3 TCP/UDP/IP High Availability Switched Configuration

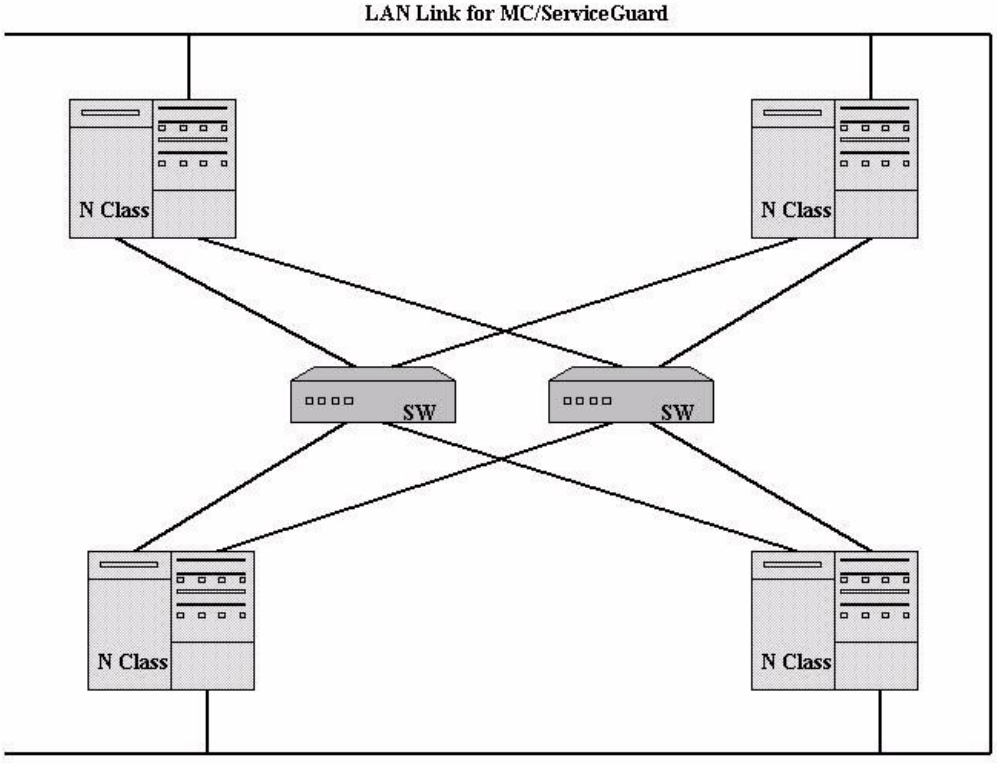

Hybrid Configuration

You can interconnect servers and workstations in a single heterogeneous HyperFabric cluster.

In this configuration, the servers are highly available. In addition, the workstations and the servers can run the same application or different applications (see Figure 2-4).

Figure 2-4 TCP/UDP/IP Hybrid Configuration

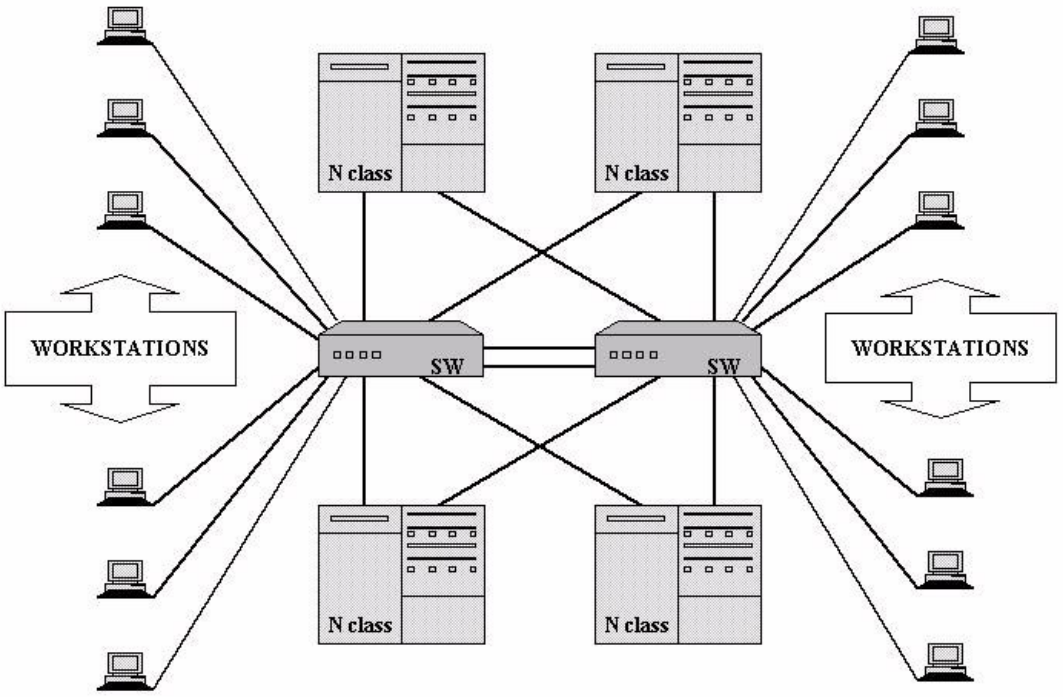

Hyper Messaging Protocol (HMP)

Hyper Messaging Protocol (HMP) is an HP patented, high performance cluster interconnect protocol. HMP provides reliable, high speed, low latency, low CPU overhead, datagram service to applications running on the HP-UX operating system.

HMP was jointly developed with Oracle Corp. The resulting feature set was tuned to enhance the scalability of the Oracle Cache Fusion clustering technology. It is implemented using Remote DMA (RDMA) paradigms.

HMP is integral to the HP-UX HyperFabric driver. It can be enabled or disabled at HyperFabric initialization using the `clic_init` command or SAM. The HMP functionality is used by the applications listed in the following “Application Availability” section.

HMP significantly enhances the performance of parallel and technical computing applications.

HMP firmware on HyperFabric adapter cards provides a shortcut that bypasses several layers in the protocol stack, boosting link performance and lowering latency. By avoiding interruptions and buffer copying in the protocol stack, communication task processing is optimized.

Application Availability

The following are the two families of applications that can use HMP over the HyperFabric interface:

- Oracle 9i Database, Release 1 (9.0.1) and Release 2 (9.2.0.1.0).

  HMP has been certified on Oracle 9i Database Release 1 with HP-UX 11.0, 11i v1, and 11i v2.

  HMP has been certified on Oracle 9i Database Release 2 with HP-UX 11.0, 11i v1, and 11i v2.

- Technical Computing Applications that use the HP Message Passing Interface (HP-MPI).

HP MPI is a native implementation of version 1.2 of the Message-Passing Interface Standard. It has become the industry standard for distributed technical applications and is supported on most technical computing platforms.

Features

The following are the HyperFabric features on HMP:

- **OnLine Addition and Replacement (OLAR)**

  The OLAR feature, which allows the replacement or addition of HyperFabric adapter cards while the system (node) is running, is supported when applications use HMP to communicate.

- **Event Monitoring Service (EMS): Supporte**d

  In the HyperFabric version B.11.23.01, the HyperFabric EMS monitor enables the system administrator to separately monitor each HyperFabric adapter on every node in the fabric, in addition to monitoring the entire HyperFabric subsystem. The monitor can inform the user if the resource being monitored is UP or DOWN. The administrator defines the condition to trigger a notification (usually a change in interface status). Notification can be accomplished with a SNMP trap, or by logging into a user specified log file with a choice of severity, or by email to a user defined email address.

  For more information on EMS, including instructions for implementing this feature, see “Configuring the HyperFabric EMS Monitor” on page 85 in this manual, and the *EMS Hardware Monitors User’s Guide* Part Number B6191-90028 September 2001 Edition.

- **ServiceGuard: Supported**

  Within a cluster, ServiceGuard groups application services (individual HP-UX processes) into packages. In the event of a single service failure (node or network), EMS provides notification and ServiceGuard transfers control of the package to another node in the cluster, allowing services to remain available with minimal interruption. ServiceGuard using EMS, directly monitors cluster nodes, LAN interfaces, and services (the individual processes within an application). ServiceGuard uses a heartbeat LAN to monitor the

nodes in a cluster. ServiceGuard cannot use the HyperFabric interconnect as a heartbeat link. Instead, a separate LAN must be used for the heartbeat.

For more information on configuring ServiceGuard, see “Configuring HyperFabric with ServiceGuard” on page 87, as well as *Managing MC/ServiceGuard* Part Number B3936-90065 March 2002 Edition.

- **High Availability (HA): Supported**

  When applications use HMP to communicate between HP 9000 nodes in a HyperFabric cluster, you can configure ServiceGuard and the EMS monitor to identify node failure and automatically failover to a functioning HP 9000 node.

  For more information on HA when running HMP applications, contact your HP representative.

- **Transparent Local Failover: Supported**

  HMP supports Transparent Local Failover in the HyperFabric version B.11.23.01.

  When a HyperFabric resource (adapter, cable, switch or switch port) fails in a cluster, HMP transparently fails over traffic using other available resources. This is accomplished using card pairs, each of which is a logical entity that comprises a pair of HF2 adapters on a HP 9000 node. Only Oracle applications can make use of the Local Failover feature. HMP traffic can only fail over between adapters that belong to the same card pair. Traffic does not fail over if both the adapters in a card pair fail. However, administrators do not need to configure HF2 adapters as card pairs if TCP/UDP/IP is run over HF2 or MPI uses HMP.

  When HMP is configured in the local failover mode, all the resources in the cluster are utilized. If a resource fails in the cluster and is restored, HMP does not utilize that resource until another resource fails.

  For more information on Transparent Local Failover while running HMP applications, see “Configuring HMP for Transparent Local Failover Support” on page 96.

- **Dynamic Resource Utilization (DRU): Partially Supported**

  If you add a new HyperFabric resource (node, cable or switch) to a cluster running an HMP application, the HyperFabric subsystem will dynamically identify the added resource and start using it. The

same process takes place when a resource is removed from a cluster. However, DRU is not supported if you add or remove an adapter from a node that is running an HMP application. This is consistent with the fact that OLAR is not supported when an HMP application is running on HyperFabric.

- **Load Balancing: Supported**

  When an HP 9000 HyperFabric cluster is running HMP applications, the HyperFabric driver balances the load across all available resources in the cluster, including nodes, adapter cards, links, and multiple links between switches.

- **Switch Management: Not Supported**

  Switch Management is not supported. Switch management will not operate properly if it is enabled on a HyperFabric cluster.

- **Diagnostics: Supported**

  You can run diagnostics to obtain information on many of the HyperFabric components using the `clic_diag, clic_probe` and `clic_stat` commands, as well as the Support Tools Manager (STM).

  For more information on HyperFabric diagnostics, see “Running Diagnostics” on page 149.

Configuration Parameters

This section discusses the maximum limits for HMP HyperFabric configurations. There are numerous variables that can impact the performance of any particular HyperFabric configuration. For more information on specific HyperFabric configurations for HMP applications, see “HMP Supported Configurations” on page 33.

- HyperFabric is supported on the HP 9000 series servers and workstations only.
- HMP is supported on the HF2 adapter, A6386A, only.
- The performance advantages that HMP offers are not completely realized unless HMP is used with A6386A HF2 (fiber) adapters and related fiber hardware. See Table 2-2 on page 20 for details. The local failover configuration of HMP is supported only on the A6386A HF2 adapters.
- Maximum Supported Nodes and Adapter Cards

HyperFabric clusters running HMP applications are limited to supporting a maximum of 64 adapter cards. However, in local failover configurations, a maximum of only 52 adapters are supported.

In point-to-point configurations running HMP applications, the complexity and performance limitations of having a large number of nodes in a cluster make it necessary to include switches in the fabric. Typically, point-to-point configurations consist of only 2 or 3 nodes.

In switched configurations running HMP applications, HyperFabric supports a maximum of 64 interconnected adapter cards.

A maximum of 8 HyperFabric adapter cards are supported per instance of the HP-UX operating system. The actual number of adapter cards a particular node is able to accommodate also depends on slot availability and system resources. See node specific documentation for details.

A maximum of 8 configured IP addresses are supported by the HyperFabric subsystem per instance of the HP-UX operating system.

- Maximum Number of Switches

  You can interconnect (mesh) up to 4 switches (16-port fiber or Mixed 8 fiber ports) in a single HyperFabric cluster.

- Trunking Between Switches (multiple connections)

  Trunking between switches can be used to increase bandwidth and cluster throughput. Trunking is also a way to eliminate a possible single point of failure. The number of trunked cables between nodes is only limited by port availability. To assess the effects of trunking on the performance of any particular HyperFabric configuration, contact your HP representative.

- Maximum Cable Lengths

  HF2 (fiber): The maximum distance is 200m (4 standard cable lengths are sold and supported: 2m, 16m, 50m and 200m).

  HMP supports up to 4 HF2 switches connected in series with a maximum cable length of 200m between the switches and 200m between switches and nodes.

  HMP supports up to 4 hybrid HF2 switches connected in series with a maximum cable length of 200m between fiber ports.

- HMP is supported on A400, A500, rp2400, rp2450, rp54xx (N-class), rp74xx (L-class), rp8400, and Superdome servers running 64-bit HP-UX.
- HMP is supported on HyperFabric starting HyperFabric versions B.11.00.11, B.11.11.01, and B.11.23.00.
- HMP is not supported on the A180 or A180C server.
- HMP is not supported on 32-bit versions of HP-UX.
- Throughput and Latency

Table 2-3 HF2 Throughput and Latency with HMP Applications

Server Class	Maximum Throughput	Latency
rp 7400	2 + 2 Gbps full duplex per link	< 22 microsec

Table 2-4 Supported HyperFabric Adapter Configurations

HF Adapter	Bus Type	Supported HP Systems	HP-UX Version	OLAR Support?	Maximum Adapters per System
A6386A	PCI (4X)	rx2600 servers	11i v2	No	1
A6386A	PCI (4X)	rx56XX servers	11i v2	No	4
A6386A	PCI (4X)	zx6000 workstations	11i v2	No	1
A6386A	PCI (4X)	SD64A servers	11i v2	Yes	8 (maximum 4 per PCI card cage)
A6386A	PCI (4X)	rx7620 servers	11i v2	No	8 (maximum 4 per PCI card cage)
A6386A	PCI (4X)	rx8620 servers	11i v2	Yes	8 (maximum 4 per PCI card cage)
A6386A	PCI (4X)	rx4640 servers	11i v2	Yes	6

NOTE The local failover configuration on HMP is supported only on the A6386A HF2 adapters.

HMP Supported Configurations

Multiple HMP HyperFabric configurations are supported to match the performance, cost and scaling requirements of each installation.

In the section, “Configuration Parameters” on page 29, the maximum limits for HMP enabled HyperFabric hardware configurations were outlined. This section discusses the HMP enabled HyperFabric configurations that HP supports. These recommended configurations offer an optimal mix of performance and availability for a variety of operating environments.

There are many variables that can impact HyperFabric performance. If you are considering a configuration that is beyond the scope of the following HP supported configurations, contact your HP representative.

Point-to-Point Configuration

You can interconnect large servers like the HP Superdome to run Oracle RAC 9i and enterprise resource planning applications. These applications are typically consolidated on large servers.

Point-to-point connections between servers support the performance benefits of HMP without investing in HyperFabric switches. This is a good solution in small configurations where the benefits of a switched HyperFabric cluster might not be required (see configuration A in Figure 2-5).

If an HMP application is running over HyperFabric and another node or adapter is added to either of the nodes, then it is necessary to also add a HyperFabric switch to the cluster (see configuration B in Figure 2-5).

Figure 2-5 HMP Point-To-Point Configurations

(Configuration A)
SINGLE CONNECTION BETWEEN TWO NODES

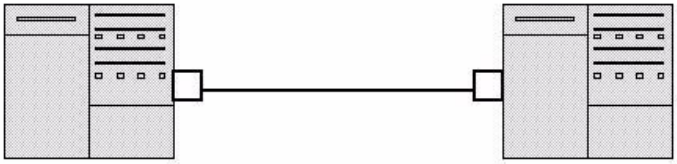

(Configuration B)
MULTIPLE CONNECTIONS BETWEEN TWO NODES

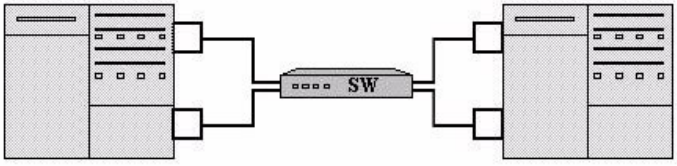

Enterprise (Database) Configuration

The HMP enterprise configuration illustrated in Figure 2-6 is very popular for running Oracle RAC 9i.

Superdomes or other large servers make up the Database Tier. Database Tier nodes communicate with each other using HMP.

Application Tier nodes communicate with each other and to the Database Tier using TCP/UDP/IP.

Figure 2-6 HMP Enterprise (Database) Configuration, Single Connection Between Nodes

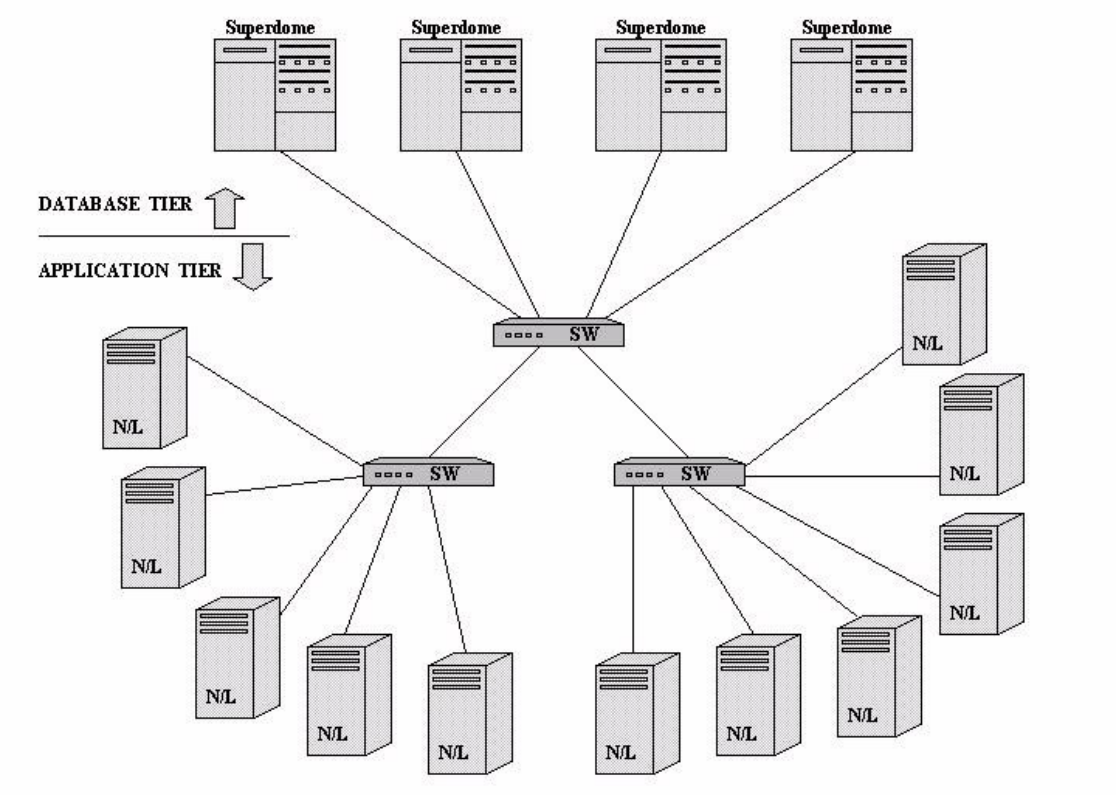

Enterprise (Database) - Local Failover Supported Configuration

The HMP enterprise configuration is a scalable solution. For high availability and performance, you can easily scale the HMP enterprise configuration with multiple connections between the HyperFabric resources. Any single point of failure in the database tier of the fabric is eliminated in Figure 2-7.

Figure 2-7 Local Failover Supported Enterprise (Database) Configuration, Multiple Connections between Nodes

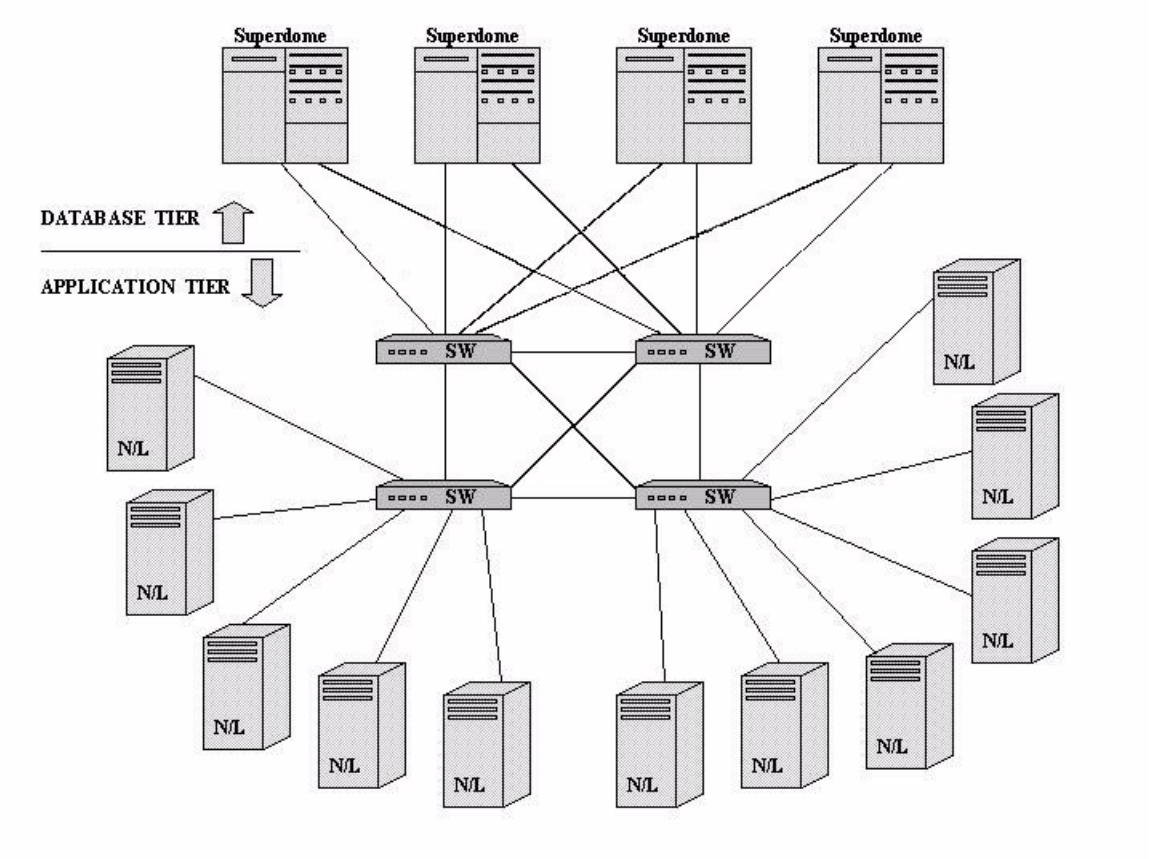

In this configuration, if a HyperFabric resource (adapter, cable, switch or switch port) fails in a cluster, HMP transparently fails over traffic using another available resource. For more information, see “Configuring HMP for Transparent Local Failover Support” on page 96.

Technical Computing (Work Stations) Configuration

This configuration is typically used to run technical computing applications with HP-MPI. A large number of small nodes are interconnected to achieve high throughput (see Figure 2-8). High availability is not usually a requirement in technical computing environments.

HMP provides the high performance, low latency path necessary for these technical computing applications. You can interconnect up to 56 nodes using HP 16-port switches. You cannot link more than four 16-port switches in a single cluster (see Figure 2-9).

The HP “J”, “B”, and “C” class workstations provide excellent performance and return on investment in technical computing configurations.

Figure 2-8 Technical Computing Configuration

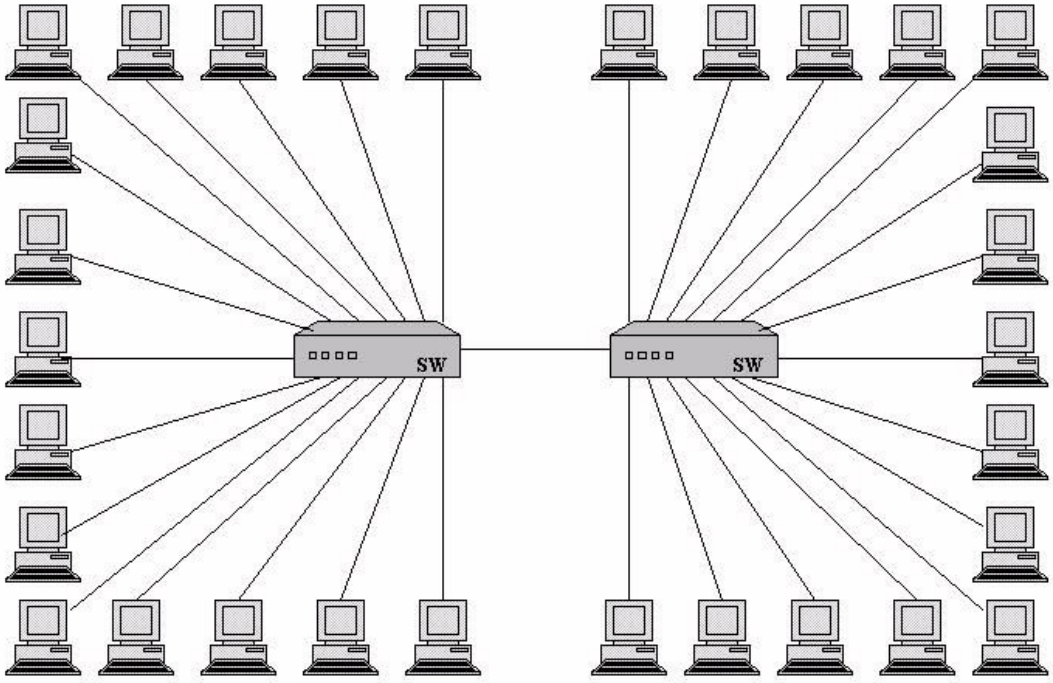

Figure 2-9 Large Technical Computing Configuration

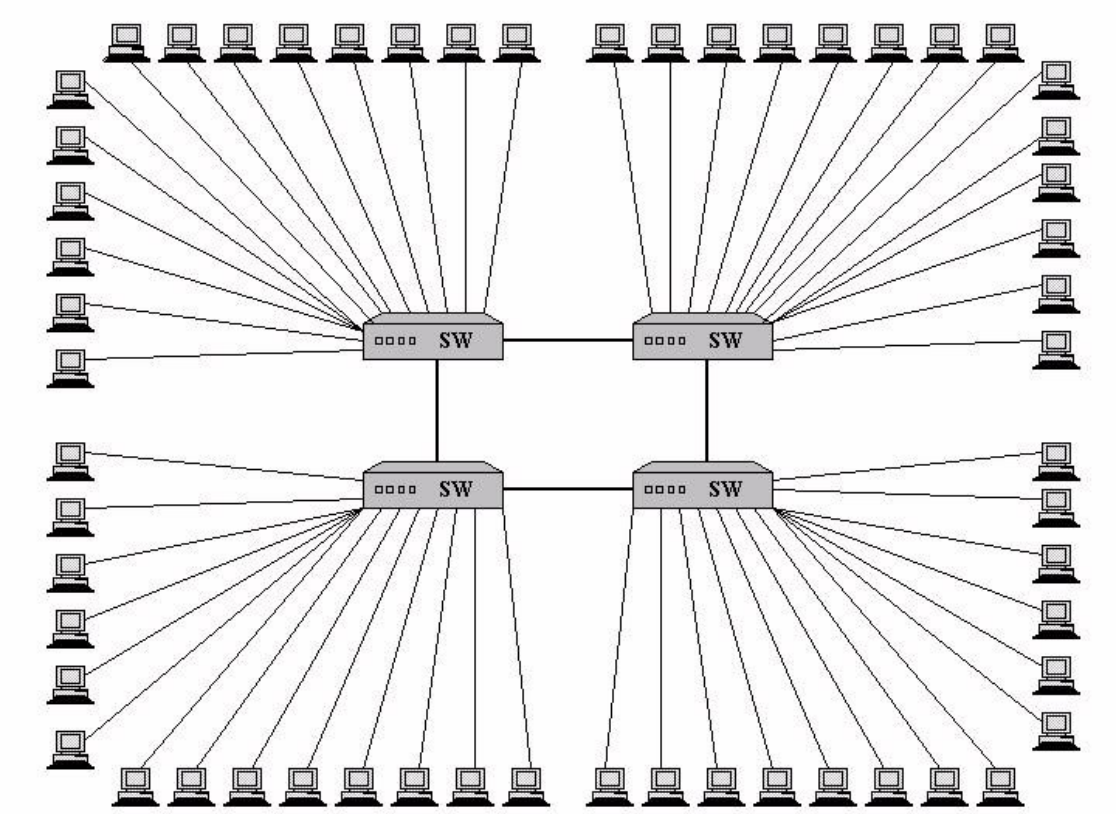

3 Installing HyperFabric

This chapter contains the following sections that describe the HyperFabric installation:

- “Installing HyperFabric Adapters” on page 44.
- “Installing the Software” on page 51.
- “Installing HyperFabric Switches” on page 57.

Checking HyperFabric Installation Prerequisites

Before installing HyperFabric, ensure that the following hardware and software prerequisites are met:

- ✓ Check the *HyperFabric Release Note*s for known problems, required patches, or other information needed for installation.
- ✓ Confirm that the `/usr/bin`, `/usr/sbin`, and `/sbin` directories are in your `PATH` by logging in as `root` and using the **`echo $PATH`** command.
- ✓ Confirm that the HP-UX operating system is the correct version. Use the **`uname -a`** command to determine the HP-UX version.

  For more information about the required operating system versions, see *HyperFabric Release Notes*.
- ✓ If you are installing an HF2 switch, confirm that you have four screws with over-sized heads.
- ✓ Confirm that there are cables of proper length and type (fiber) to make connections in the fabric (adapter to adapter, adapter to switch, or switch to switch).
- ✓ Confirm that there is at least one loopback plug for testing the adapters and switches (a fiber loopback plug [HP part number A6384-67004] is shipped with each HF2 switch).
- ✓ Confirm that necessary tools are available to install the HyperFabric switch mounting hardware. In addition, check the HP 9000 system's documentation to determine if additional tools may be required for component installation.
- ✓ Confirm that the software media is correct.
- ✓ Create a map of the fabric (optional).
- ✓ Confirm that HP-UX super-user privileges are available; they are necessary to complete the HyperFabric installation.

The first HyperFabric installation step is installing HyperFabric adapter cards in the nodes.

Installing HyperFabric Adapters

This section contains information about installing HyperFabric adapters in HP 9000 systems. Online Addition and Replacement (OLAR) information is provided in the section, “Online Addition and Replacement” on page 44.

CAUTION

HyperFabric adapters contain electronic components that can easily be damaged by small amount of electricity. To avoid damage, follow these guidelines:

- Store adapters in their antistatic plastic bags until installation.
- Work in a static-free area, if possible.
- Handle adapters by the edges only. Do not touch electronic components or electrical traces.
- Use the disposable grounding wrist strap provided with each adapter. Follow the instructions included with the grounding strap.
- Use a suitable ground—any exposed metal surface on the computer chassis.

For specific instructions, see system-specific documentation on “installing networking adapters” for each type of HP 9000 system that HyperFabric adapters will be installed into.

When the HyperFabric adapters have been installed, go to “Installing the Software” on page 51.

Online Addition and Replacement

Online Addition and Replacement (**OLAR**) allows PCI I/O cards, adapters or controllers to be replaced or added to HP 9000 systems, without the need for completely shutting down and rebooting the system, or adversely affecting other system components. This feature is only

available on HP 9000 systems that are designed to support OLAR. The system hardware uses the per-slot power control combined with OS support to enable this feature.

NOTE OLAR is supported only on TCP/UDP/IP over HF2 adapters.

Not all add-in cards have this capability, but over time many cards will be gaining this capability.

The *HyperFabric Release Note*s contains information about which HP 9000 systems and HyperFabric adapters OLAR is supported for.

IMPORTANT At this time, Superdome systems are not intended for access by users. HP recommends that these systems only be opened by a qualified HP engineer. Failure to observe this requirement can invalidate any support agreement or warranty to which the owner might otherwise be entitled.

There are two methods to add or replace OLAR-compatible cards:

- Using the SAM utility.
- Issuing command-line commands, through `olrad`, that refer to the HyperFabric OLAR script (`/usr/sbin/olard.d/clicd`).

HP recommends that SAM be used for OLAR procedures, instead of the `rad` command. This is primarily because SAM prevents the user from doing things that might have adverse effects. This is not true when the `rad` command is used.

For detailed information about using either of these two procedures, see *Configuring HP-UX For Peripherals*. You can order that document from HP, or you can view, download, and print it from the following URL: http://www.docs.hp.com.

Table 3-1 below explains some important OLAR-related terms.

Table 3-1 Important OLAR Terms

Term	Meaning
OLAR	All aspects of the OLAR feature including Online Addition (OLA) and Online Replacement (OLR).
Power Domain	A grouping of 1 or more interface card slots that are powered on or off as a unit. (**Note**: Multi-slot power domains are not currently supported.)
target card / target card slot	The interface card which will be added or replaced using OLAR, and the card slot in which it resides.
affected card / affected card slot	Interface cards and the card slots they reside in, which are in the same power domain as the target slot.

IMPORTANT In many cases, other interface cards and slots within the system are dependent on the target card. For example, if the target card is a multiple-port card, suspending or deleting drivers for the target card slot also suspends individual drivers for the multiple hardware paths on that card.

During a card replacement operation, SAM performs a **Critical Resource Analysis** (**CRA**), which checks all ports on the target card for critical resources that would be temporarily unavailable while the card is shut down.

Planning and Preparation

As mentioned previously, for the most part, SAM prevents the user from performing OLAR procedures that would adversely affect other areas of the HP 9000 system. See *Configuring HP-UX For Peripherals* for detailed information.

Critical Resources

The effects of shutting down a card's functions must be considered. Replacing a card that is still operating can have extensive consequences. Power to a slot must be turned off when a card is removed and a new card is inserted.

This is particularly important if there is no online failover or backup card to pick up those functions. For example:

- Which mass storage devices will be temporarily disconnected when a card is shut down?
- Will a critical networking connection be lost?

A critical resource is one that would cause a system crash or prevent an operation from successfully completing if the resource were temporarily suspended or disconnected. For example, if the SCSI controller is connected to the unmirrored root disk or swap space, the system will crash when the SCSI controller is shut down.

During an OLAR procedure, it is essential to check the targeted card for critical resources, as well as the effects of existing disk mirrors and other situations where a card's functions can be taken over by another card that will not be affected.

As mentioned earlier, SAM performs a thorough CRA automatically, and presents options based on its findings. If it is determined that critical resources will be affected by the OLAR procedure, the card could be replaced when the system is offline. If action must be taken immediately, an online addition of a backup card and deletion of the target card could be attempted using the rad command.

Card Compatibility

This section explains card compatibility considerations for doing OLAR.

Online Addition (OLA) Multiple cards can be added at the same time. When adding a card online, the first issue to resolve is whether the new card is compatible with the system. Each OLAR-capable PCI slot provides a set amount of power. The replacement card cannot require more power than there is available.

The card must also operate at the slot's bus frequency. A PCI card must run at any frequency lower than its maximum capability, but a card that could operate at only 33 MHz would not work on a bus running at 66 MHz. `rad` provides information about the bus frequency and power available at a slot, as well as other slot-related data.

If an HP 9000 system has one or more slots that support OLAR and OLA will be used to install a HyperFabric adapter in one of those slots—install the adapter in the HP 9000 system according to the procedure described in the "Managing PCI Cards with OLAR" chapter of the *"Configuring HP-UX Peripherals"* manual.

After adding a new HyperFabric adapter, SAM tries to locate the HyperFabric software. If SAM cannot locate the HyperFabric software, the new adapter cannot be used until the software is installed (software installation requires a system reboot). If SAM locates the HyperFabric software, SAM determines whether the new adapter is functional. If it is not functional, SAM displays an error message.

If the new adapter is functional, SAM displays a message telling the user to configure the adapter and start HyperFabric. If only one adapter is being added, issue the `clic_init -c` command or use SAM to configure the adapter, and then issue the `clic_start` command or use SAM to start HyperFabric. If multiple adapters are being added, add all of the adapters first, and then run `clic_init -c` and `clic_start` or use SAM. See "Performing the Configuration" on page 78 and "Starting HyperFabric" on page 109 for more information about configuring and starting HyperFabric.

CAUTION Do not change any configuration information for an existing HyperFabric adapter or switch while you are using `clic_init -c` to configure a new adapter.

When you have completed the adapter installation, go to "Installing the Software" on page 51.

Online Replacement (OLR) When replacing an interface card online, the replacement card must be identical to the card being replaced (or at least be able to operate using the same driver as the replaced card). This is referred to as **like-for-like** replacement and should be adhered to, because using a similar but not identical card can cause unpredictable results. For example, a newer version of the target card that is identical to the older card in terms of hardware might contain an updated firmware version that could potentially conflict with the current driver. An A6386A adapter must be replaced with another A6386A adapter, etc. Also, the old adapter and new adapter must have the same revision levels.

When a replacement card is added to an HP 9000 system, the appropriate driver for that card must be configured in the kernel before beginning the replacement operation. SAM ensures the correct driver is present. (In most cases, the replacement card will be the same type as a card already in the system, and this requirement will be automatically met.) Keep the following things in mind:

- If the necessary driver is not present and the driver is a dynamically loadable kernel module (DLKM), it can be loaded manually. See the "Dynamically Loadable Kernel Modules" section in *"Configuring HP-UX For Peripherals"* for more information.
- If the driver is static and not configured in the kernel, then the card cannot be added online. The card could be physically inserted online, but no driver would claim it.

If there is any question about the driver's presence, or if it is uncertain that the replacement card is identical to the existing card, `ioscan` can be used together with `rad` to investigate.

If more than one operational HyperFabric adapter is present when SAM requests the suspend operation for all ports on the target adapter, HyperFabric will redirect the target adapter's traffic to a local backup adapter using local failover. Client applications using the replaced adapter will not be interrupted in any way.

If the adapter being replacing is active and it is the only operational HyperFabric adapter on the HP 9000 system, SAM displays the following warning message:

```
WARNING: You have 1 operational HyperFabric card. If you go
ahead with this operation you will lose network access via
HyperFabric until the on-line replaced HyperFabric card
becomes operational.
```

You are asked if you want to continue. If you reply Yes, client applications are suspended. Replace the adapter according to the procedure described in the “Managing PCI Cards with OLAR” chapter of the *Configuring HP-UX Peripherals* manual.

When an adapter has been replaced, client application activity resumes unless the TCP timers or the application timers have popped.

CAUTION

Do not use the `clic_start` command or the `clic_shutdown` command, while an installed adapter is suspended. Do not use SAM to start or stop HyperFabric while an installed adapter is suspended. The operation will fail and an error message will be displayed.

After a HyperFabric adapter has been replaced, SAM checks the replacement adapter to ensure that it is permitted according to the like-for-like rules. If the adapter is permitted, SAM automatically activates it. If it is not permitted, SAM displays an error message.

Installing the Software

This section describes the HyperFabric file structure and the steps necessary to load the software. The software must be installed on each instance of the HP-UX operating system in the fabric.

File Structure

The HyperFabric file structure is shown in Figure 3-1 below. The structure is shown for informational purposes only. The user cannot modify any of the files or move them to a different directory.

Figure 3-1 HyperFabric File Structure

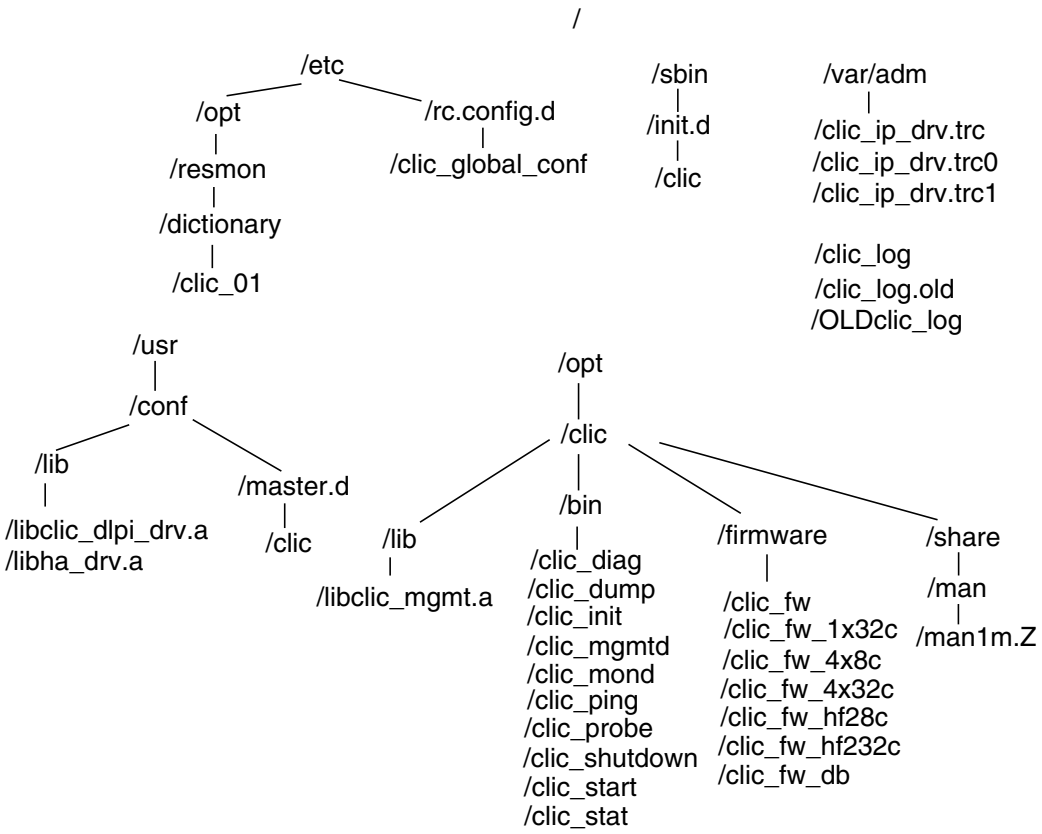

The commands and files used to administer HyperFabric typically have a prefix of `clic_`. CLIC stands for CLuster InterConnect, and it is used to differentiate those HyperFabric commands/files from other commands/files. For example, the HyperFabric command `clic_init` is different from the HP-UX `init` command.

Each of the files shown in Figure 3-1 above is briefly described below:

- `/etc/opt/resmon/dictionary/clic_01`

  The HyperFabric dictionary file for the Event Monitoring Service (EMS).

- `/etc/rc.config.d/clic_global_conf`

  The global configuration file, which contains the IP addresses for each adapter and each HyperFabric switch (if any) in the fabric.

- `/sbin/init.d/clic`

  The system boot startup script for the HyperFabric management process.

- `/var/adm/clic_ip_drv.trc`

  One of the software's trace files. This file is created when the `clic_diag -D TCP_IP` command is run.

- `/var/adm/clic_ip_drv.trc0`

  One of the HyperFabric software's trace files. This is the primary file that is created when the `clic_diag -C TCP_IP` command is run.

- `/var/adm/clic_ip_drv.trc1`

  One of the HyperFabric software's trace files. This file is created when the `clic_diag -C TCP_IP` command is run, and the primary trace file (`clic_ip_drv.trc0`) becomes full.

- `/var/adm/clic_log`

  The global log file that is updated by the HyperFabric management process.

- `/var/adm/clic_log.old`

  The backup copy of the log file that is created when the log file grows larger than 100 Kbytes.

- `/var/adm/OLDclic_log`

  The log file from the previous time the `clic_start` command was executed.

- `/usr/conf/lib/libclic_dlpi_drv.a`

  The kernel library that contains the HyperFabric software.

- `/usr/conf/lib/libha_drv.a`

  The kernel library that contains the High Availability (HA) software.

- `/usr/conf/master.d/clic`

  This file is described along with the other master files in the `master` man page (type `man master` at the HP-UX prompt).

- `/opt/clic/lib/libclic_mgmt.a`

  The HyperFabric management API library.

- `/opt/clic/bin`

  The directory containing the HyperFabric management commands: `clic_diag`, `clic_init`, `clic_probe`, `clic_shutdown`, `clic_start`, `clic_stat`, and `clic_dump`. (Note that `clic_dump` is for HP internal use only.) Also, `clic_ping` is replaced by `clic_probe`. This directory also contains the HyperFabric management process (`clic_mgmtd`) and the HyperFabric EMS monitor process (`clic_mond`).

- `/opt/clic/firmware/clic_fw`

  The 1X HSC HyperFabric 8-bit CRC firmware. This file must not be modified.

- `/opt/clic/firmware/clic_fw_1x32c`

  The 1X HSC HyperFabric 32-bit CRC firmware. This file must not be modified.

- `/opt/clic/firmware/clic_fw_4x8c`

  The 4X PCI HyperFabric 8-bit CRC firmware. This file must not be modified.

- `/opt/clic/firmware/clic_fw_4x32c`

  The 4X HyperFabric PCI 32-bit CRC firmware. This file must not be modified.

- `/opt/clic/firmware/clic_fw_hf28c`

  The HyperFabric2 8-bit firmware. This file must not be modified.

- `/opt/clic/firmware/clic_fw_hf232c`

  The HyperFabric2 32-bit firmware. This file must not be modified.

- `/opt/clic/firmware/clic_fw_db`

  A binary file where adapter-specific configuration information is stored. The management process creates this file using default values.

- `/opt/clic/share/man/man1m.Z`

  The man pages for the HyperFabric commands.

Loading the Software

Listed below are the steps you must follow to load the HyperFabric software, using the HP-UX `swinstall` program.

Step 1. Log on to the system as `root`.

Step 2. Insert the software media into the appropriate drive. If the software is being loaded from a CD-ROM, go to step 3; otherwise, go to step 4.

Step 3. Mount the CD-ROM drive by using the following command at the command prompt:

```
$ mount device_name
```

where *device_name* is the name assigned to the CD-ROM drive.

Step 4. Run the `swinstall` program using the following command:

```
$ /usr/sbin/swinstall
```

This opens the “Software Selection” window.

Step 5. Change the `Source Host Name`, if necessary, and then enter the mount point of the drive in the `Source Depot Path` field. Select the `OK` button to return to the “Software Selection” window.

The “Software Selection” window now contains a list of available software to install.

Step 6. Highlight the HyperFabric software:

- HP-UX 11i v2: `HyprFabrc-00`

Step 7. Choose `Mark for Install` from the “Actions” menu; this chooses the highlighted software.

Step 8. From the “Actions” menu, select the “Install...” menu, and then choose `Install`. This begins product installation and opens the “Install Analysis” window.

Step 9. Select the `OK` button in the “Install Analysis” window when the `Status` field displays a “Ready” message.

Step 10. Select the `YES` button in the “Confirmation” window to start software installation.

`swinstall` loads the fileset, runs the control script for the filesets, and builds the kernel. When the processing is finished, the "Status" field displays a "Ready" message. Select "Done" and then the "Note" window opens.

Step 11. Select the `OK` button in the "Note" window to reboot. The user interface disappears and the system reboots.

Step 12. When the system comes back up, log on to the system as `root` and view the `/var/adm/sw/swagent.log` and `/var/adm/sw/swinstall.log` files to view error or warning messages that may have occurred during the installation.

Step 13. While still logged in as `root`, view the `/etc/services` file to ensure that these two HyperFabric-related lines are present:

- `hp-clic 3384/tcp #clic management daemon`
- `hp-clic 3384/udp #clic switch management`

These lines are used by the HyperFabric software, and are not comments; therefore, do not remove them from the file.

Step 14. Verify that all installed HyperFabric adapters have a software state of "CLAIMED," by running the **`ioscan -nf -C clic`** command.

A check is also done to make sure all of the HyperFabric adapters have been claimed when `clic_init` is activated or when SAM is used to configure HyperFabric.

Step 15. If one or more HyperFabric switches are included in the configuration, go to the next section of this chapter, "Installing HyperFabric Switches"; otherwise, go to Chapter 4, "Configuring HyperFabric," on page 67.

Installing HyperFabric Switches

This section contains the information you need to install HyperFabric switches. As stated earlier, the term **HyperFabric2 (HF2) switch** refers to the functional switch (the A6384A switch chassis with one of the switch modules installed).

Before Installation

Before you install the HyperFabric switch, you should be aware of the following:

- The A6384A HF2 switch is supported beginning with the following HyperFabric software versions:
  - HP-UX 11.0: version B.11.00.11
  - HP-UX 11i v1: version B.11.11.01
  - HP-UX 11i v2: version B.11.23.00.

  HyperFabric switches are not supported by software versions earlier than those mentioned above, respectively.

  To determine the version of HyperFabric you have, issue the following command at the command prompt:

  ```
  $ swlist | grep -i hyperfabric
  ```

- For the HF2 switch, HP recommends that you use the rails shipped with the switch when you mount it in a standard 19-inch rack, even though the switch can be mounted in the rack by itself (without the rails).

WARNING **To prevent overheating, you must leave one rack unit (1 EIA) of empty space above the HyperFabric switch.**

- After the HyperFabric switch is mounted in the rack, attach the various cables to the switch.

  To avoid damage to any of the cables, follow these guidelines:

— If your cables have dust caps over the connectors, keep them in place until you are ready to connect them. This prevents dirt and oils from soiling any important surfaces.

— Be careful not to stretch, puncture, or crush the cable.

To install an HF2 switch, see “Installing the HF2 Switch” on page 59.

Installing the HF2 Switch

This section contains information on installing an HF2 switch.

The front of the HF2 switch has a flange—or "wing"—on each side, with two holes for attaching the switch to the rack. The following figures do not show the flanges.

Figure 3-2 shows the front of the HF2 switch with an A6388A HF2 8-port fiber switch module installed in the switch's expansion slot.

Figure 3-2 Front of HF2 Switch (A6388A Switch Module Installed)

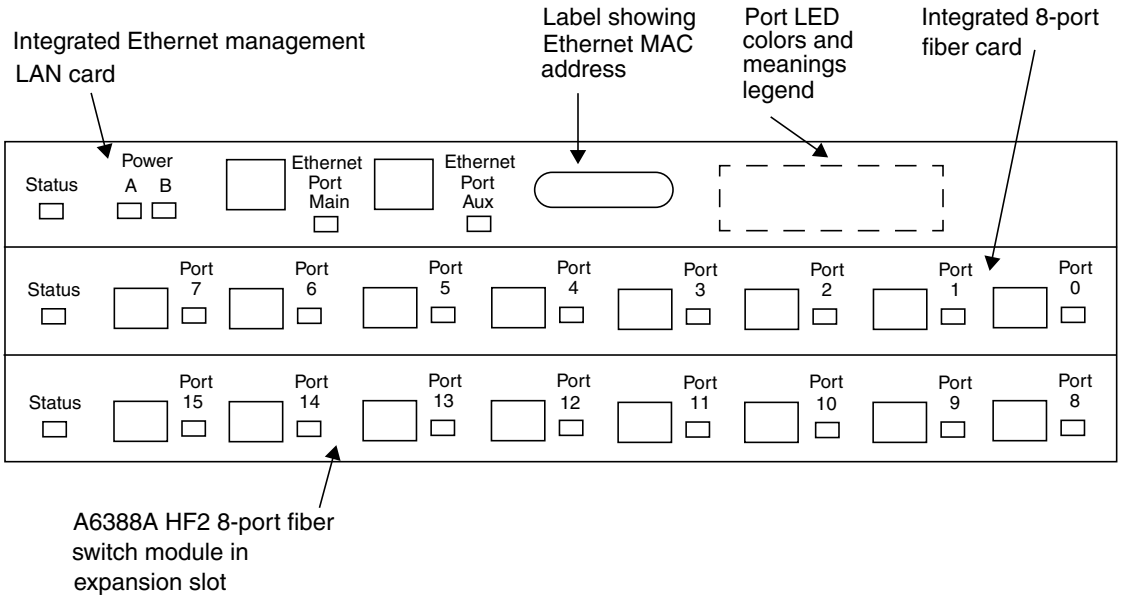

Figure 3-3 shows the front of the HF2 switch with an A6389A HF2 4-port copper switch module installed in the switch's expansion slot.

Figure 3-3 Front of HF2 Switch (A6389A Switch Module Installed)

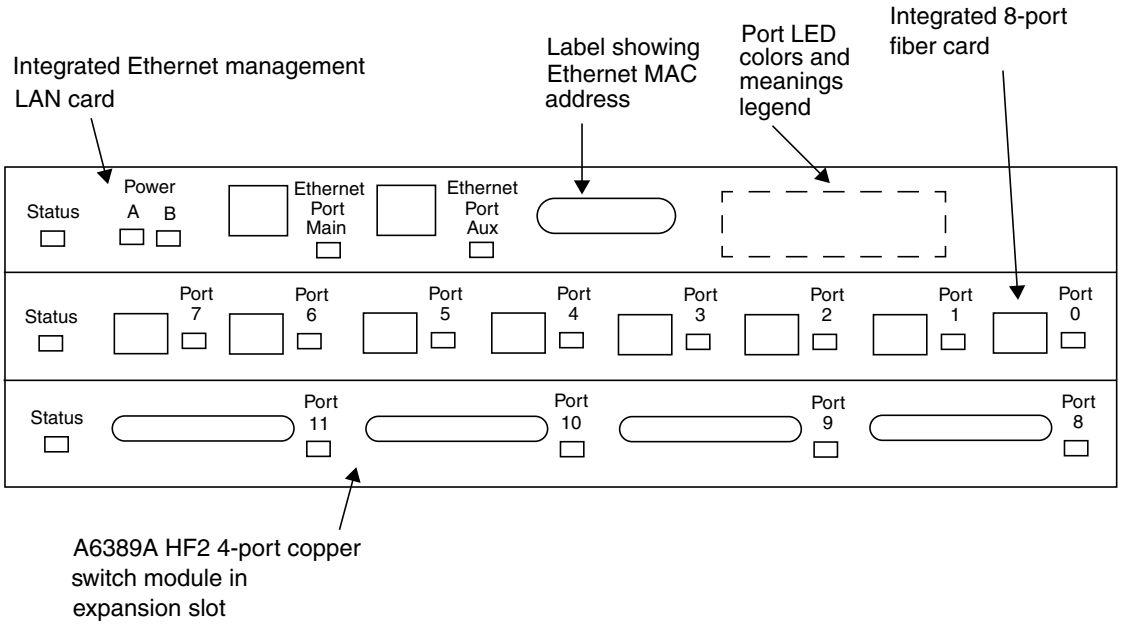

You can install the HF2 switch in one of the following ways:

- Using the rail kit that is shipped with the switch (see the next section, "With the Rail Kit"). HP strongly recommends installing the HF2 switch this way.
- Attaching the switch directly to the rack (see "Without the Rail Kit" on page 65).

With the Rail Kit

As mentioned earlier, HP strongly recommends installing the HF2 switch using the rail kit.

When you install the HF2 switch, you will be putting the front of the switch—the end with the flanges (“wings”)—at the back of the rack. The steps for installing the HF2 switch using the rail kit are as follows:

Step 1. Prepare the rack for rail and switch installation.

Step 2. Install and secure the rails in the rack, using two screws per rail. The following figure shows the rack with the rails installed.

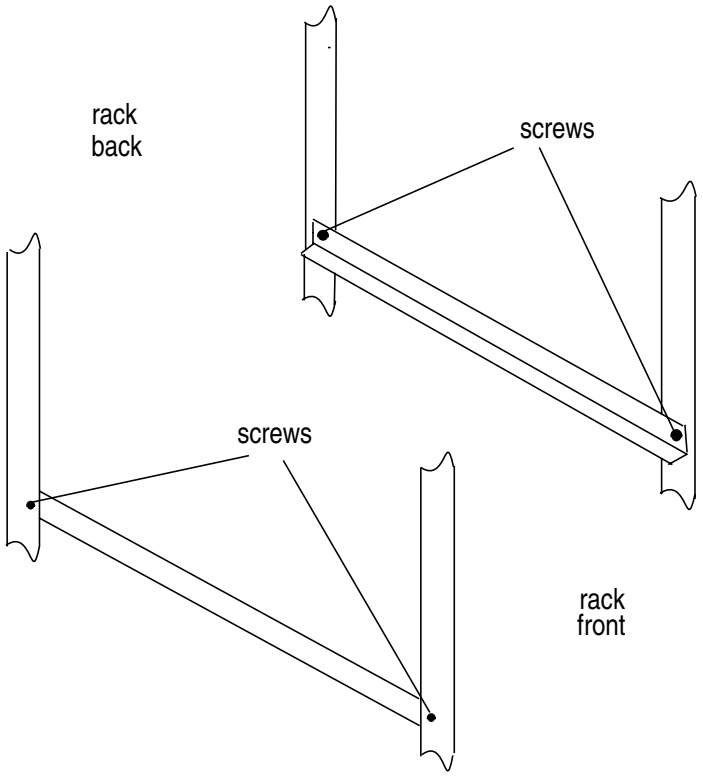

Step 3. From the front of the rack, slide the switch—with the front of the switch facing the back of the rack—into the rack, on the rails. Move it until it is snug against the back of the rack. You might not have enough clearance between the switch and the rail screws; this may prevent you from easily

sliding the switch. If so, try lifting the switch over the rail screws. If you cannot do this, remove the rail screws, slide the switch into position, and put the rail screws back in.

Step 4. Align the two holes in each flange ("wing") on the switch's front with the holes in the rack frame. Fasten each flange of the switch to the rack by putting a screw in each of the four holes in the flanges. Be sure to use screws with over-sized heads. The following figure shows the rack with the switch in this position.

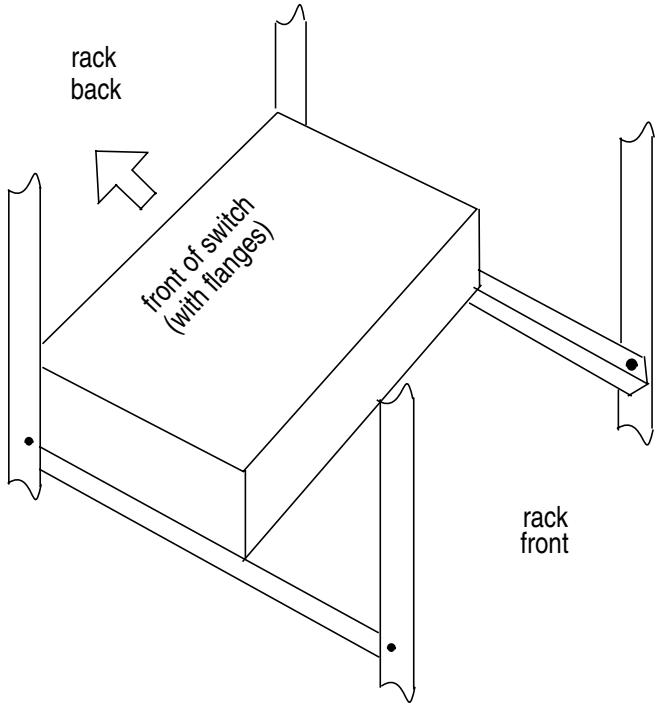

Step 5. From the front of the rack, install a bracket on the outside of each rail, using two screws per bracket. Be sure to use the upper screw holes on each bracket. Put the screws in the two square holes closest to the back of the switch, so that the brackets—referred to as "bracket 1" and "bracket 2" in these steps—are snug against the switch. Be sure that the brackets make contact with the beveled (sloped) part of the switch (not

the flat part). Otherwise, the brackets will not secure the back of the switch properly. The following figure shows the rack with these two brackets installed.

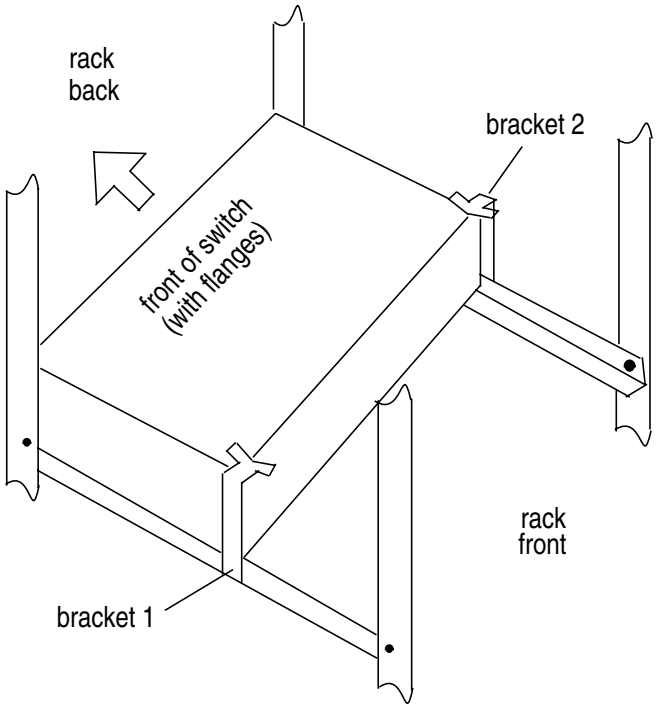

Step 6. When you are sure the brackets are snug against the beveled part of the switch, tighten all four screws in brackets 1 and 2.

Step 7. For each port that will be connected to an HyperFabric adapter in an HP 9000 system, attach the cable from the corresponding adapter. Your connections must be copper-to-copper and fiber-to-fiber (including cables).

Step 8. Connect the switch to the Ethernet network.

Step 9. Plug the switch's power cord into the rack's PDU, if it has one. Alternatively, you can plug a power cord that is compatible with your country's requirements into a power strip or outlet that you want to use for the switch. (In this case, you are responsible for obtaining a compatible power cord.)

Step 10. Power on the HF2 switch by plugging the power cord into the AC inlet on the back of the switch. (There is no power switch.)

Step 11. Once the power is on, check these LEDs on the integrated Ethernet management LAN adapter card (in the top slot of the switch):

- ✓ The "Operating/Fault" LED shows solid green.
- ✓ The "Power A" and "Power B" LEDs show solid green.
- ✓ The "Ethernet Port Main" and "Ethernet Port Aux" LEDs are showing solid green (connected) or flashing green (Ethernet traffic is flowing to the switch). See Figure 3-2 on page 59 or Figure 3-3 on page 60 for the locations of the LEDs.

Step 12. On the integrated 8-port fiber card (in the middle slot of the switch), check that for each switch port that is connected to an HF2 adapter, the LED on the port shows as solid green (see Figure 3-2 on page 59 or Figure 3-3 on page 60). This means the connection is operational.

Step 13. On the switch module in the expansion slot (the bottom slot of the switch), check that for each switch port that is connected to an HyperFabric adapter, the LED on the port shows as solid green (see Figure 3-2 on page 59 or Figure 3-3 on page 60). This means the connection is operational.

For more information about the switch's LEDs, see "HF2 Switch LEDs" on page 148.

Step 14. If you want to install another HF2 switch using the rail kit, go back to step 1.

If you want to install another HF2 switch without using the rail kit, go to the next section, "Without the Rail Kit".

Otherwise, go to Chapter 4, "Configuring HyperFabric," on page 67.

Without the Rail Kit

As mentioned earlier, HP strongly recommends installing the HF2 switch using the rail kit (described in the previous section, “With the Rail Kit” on page 61).

When you install the HF2 switch, you will be putting the front of the switch—the end with the flanges (“wings”) — at the back of the rack. The steps for installing the HF2 switch without using the rail kit are as follows:

Step 1. Prepare the rack for switch installation.

Step 2. Insert the HF2 switch into the rack, with the front of the switch snug against the back of the rack.

Step 3. Align the two holes in each flange on the switch’s front with the holes in the rack frame.

Step 4. Fasten each flange of the switch to the rack by putting a screw in each of the four holes in the flanges. Be sure to use screws with over-sized heads.

Step 5. Tighten all of the screws so that the HF2 switch is firmly mounted in the rack.

Step 6. For each port that will be connected to an HyperFabric adapter in an HP 9000 system, attach the cable from the corresponding adapter. Your connections must be copper-to-copper and fiber-to-fiber (including cables).

Step 7. Connect the switch to the Ethernet network.

Step 8. Plug the switch’s power cord into the rack’s PDU, if it has one. Alternatively, you can plug a power cord that is compatible with your country’s requirements into a power strip or outlet that you want to use for the switch. (In this case, you are responsible for obtaining a compatible power cord.)

Step 9. Power on the HF2 switch by plugging the power cord into the AC inlet on the back of the switch. (There is no power switch.)

Step 10. Once the power is on, check these LEDs on the integrated Ethernet management LAN adapter card (in the top slot of the switch):

- ✓ The “Operating/Fault” LED shows solid green.
- ✓ The “Power A” and “Power B” LEDs show solid green.
- ✓ The “Ethernet Port Main” and “Ethernet Port Aux” LEDs are showing solid green (connected) or flashing green (Ethernet traffic is flowing to the switch). See Figure 3-2 or Figure 3-3 below for the locations of the LEDs.

Step 11. On the integrated 8-port fiber card (in the middle slot of the switch), check that for each switch port that is connected to an HF2 adapter, the LED on the port shows as solid green (see Figure 3-2 on page 59 or Figure 3-3 on page 60). This means the connection is operational.

Step 12. On the switch module in the expansion slot (the bottom slot of the switch), check that for each switch port that is connected to a HyperFabric adapter, the LED on the port shows as solid green (see Figure 3-2 on page 59 or Figure 3-3 on page 60). This means the connection is operational.

For more information about the switch’s LEDs, see “HF2 Switch LEDs” on page 148.

Step 13. If you want to install another HF2 switch without using the rail kit, go to step 1.

If you want to install another HF2 switch using the rail kit, go to “With the Rail Kit” on page 61.

Otherwise, go to Chapter 4, “Configuring HyperFabric,” on page 67.

4 Configuring HyperFabric

This chapter contains the following sections that describe configuring HyperFabric:

Configuration Overview

You do not need to configure the HyperFabric switch because the HyperFabric management process performs automatic routing and configuring for the switch. So, configuring HyperFabric consists only of creating the HyperFabric `/etc/rc.config.d/clic_global_conf` global configuration file on each node in the fabric. The configuration file contains the following information:

- The IP addresses and subnet mask of the HyperFabric adapters installed in the node.
- For each HyperFabric switch in the fabric — the switch's IP address, and the MAC address of the switch's Ethernet port. This applies only if you enable switch management. In addition, you cannot enable switch management through SAM — you must use the `clic_init` command.
- The IP multicast address that all the switches and nodes in the fabric will register to (if you are going to enable switch management).
- The IP address of the local node's Ethernet LAN interface. This LAN interface must be on the same subnet as Ethernet ports of the HyperFabric switches (if you are going to enable switch management). (A node might have multiple LAN interfaces.)

NOTE HP recommends that you do not enable switch management.

You can create the global configuration file by either running the `clic_init` command or using SAM to configure each HyperFabric adapter.

The `clic_init` command and SAM also place the necessary entries into the following three files:

- The system `/etc/rc.config.d/netconf` file.

IMPORTANT

In this file, `clic_init` and SAM add some HyperFabric-related lines that end with the characters `#clic`. These lines are used by the HyperFabric software — and are not comments — so do not remove them from the file.

- The system `/etc/rc.config.d/clic_global_conf` file.
- The `/etc/rarpd.conf` (Reverse Address Resolution Protocol [RARP]) support file. This file is used in the management of the HyperFabric switches (if you are going to enable switch management).

The `clic_init` command is described in “Using the clic_init Command” on page 79. Using SAM to configure an adapter is described in “Using SAM” on page 81.

After you have used the `clic_init` command or SAM, you can configure HyperFabric with ServiceGuard, if necessary. For more information, see “Configuring HyperFabric with ServiceGuard” on page 87.

You can configure card pairs for the Transparent Local Failover Feature of HMP available with the B.11.23.01 release of HyperFabric. For more information, see “Configuring HMP for Transparent Local Failover Support” on page 96.

Information You Need

When you run the `clic_init` command or use SAM for configuration, you have to provide certain configuration information. So, before you run `clic_init` or use SAM, you should have the following information:

- ❏ For each node in the fabric, determine if that node will need to interoperate with other nodes that are using; any HP-UX 11.0 HyperFabric versions earlier than B.11.00.11 or any HP-UX 11i v1 HyperFabric versions earlier than B.11.11.01.
- ❏ For each HyperFabric adapter installed in the local node:
  - ✓ The adapter's IP address.

IMPORTANT

The last 10 bits of each adapter's IP address must be unique throughout the entire fabric. And, remember that the last part of the address cannot be `0` (that is, the IP address cannot be *n.n.n.*`0`). Also, note that HyperFabric converts these 10 bits to a decimal value called the **Virtual Route IDentifier** (**VRID**), which is used in some HyperFabric command input and output.

  - ✓ The subnet mask. When you run `clic_init` or use SAM, if you do not specify a value for this, a default subnet mask is chosen based on the adapter's IP address.

  When `clic_init` begins to prompt you for the information for each adapter, it assigns an ID (for example, `clic0`) to that adapter and displays it as part of the first prompt. If you use SAM, it assigns the adapter an ID and displays it in the "Adapter Name" column of the "Configure HyperFabric Adapter" screen. You can also determine an adapter's ID by running the `clic_stat` command (see "The clic_stat Command" on page 117). You should note each adapter's ID, because it is used as input to other HyperFabric commands.

- ❏ For using the Transparent Local Failover feature of HMP available in the version B.11.23.01 of HyperFabric, you need to define the card pairs.
- ❏ For each HyperFabric switch in the fabric (if you are going to enable switch management):

- ✓ The IP address of the switch.
- ✓ The MAC address of the switch's Ethernet port. If you do not already know the switch's MAC address, it is printed on a label on the back of the HF switch and on the front of the HF2 switch. See Figure 3-2 on page 66 for the location of the label on the HF switch, and Figure 3-2 on page 59 and Figure 3-3 on page 60 for the location of the label on the HF2 switch.

IMPORTANT

You cannot enable switch management through SAM — you must use the `clic_init` command.

When `clic_init` begins to prompt you for the information for each switch, it assigns an ID (for example, `sw_clic0`) to that switch and displays it as part of the first prompt. Note that you can also determine a switch's ID by running the `clic_stat` command (see "The clic_stat Command" on page 117). You should note each switch's ID, because it is used as input to other HyperFabric commands.

- ❏ For the entire fabric, you need the IP multicast address that all the switches and nodes in the fabric will register to. The address must be a class D address. Note that if you do not have switch management enabled, you do not need this information (`clic_init` will not prompt you for it).
- ❏ For each node in the fabric, you need the IP address of the node's Ethernet LAN interface that is on the same subnet as the switches. (As mentioned earlier, a node might have multiple LAN interfaces.) Note that if you do not have switch management enabled, you do not need this information (`clic_init` will not prompt you for it).

As stated earlier, HP recommends that you do not enable switch management.

IMPORTANT

You should also check your `/etc/hosts` file — when you are using files for host name look up — to ensure that the entries for all of the systems are in the correct format: the official host name, which is the full domain extended host name, and any alias names. For example:

```
IP_address   bently6.corp3.com   bently6
IP_address   bently4.corp7.com   test1
IP_address   bently2.corp4.com   test3
```

Configuration Information Example

This example uses some “dummy” (that is, not valid) addresses to the components in Figure 4-1. The dummy addresses are used only to show the flow of the information provided as input to the `clic_init` command and SAM. Do not try to use these addresses in your configuration.

To use `clic_init` to configure the Transparent Local Failover feature on HMP, see the section, “Configuring HMP for Transparent Local Failover Support - Using the clic_init command” on page 104.

Figure 4-1 Map for Configuration Information Example

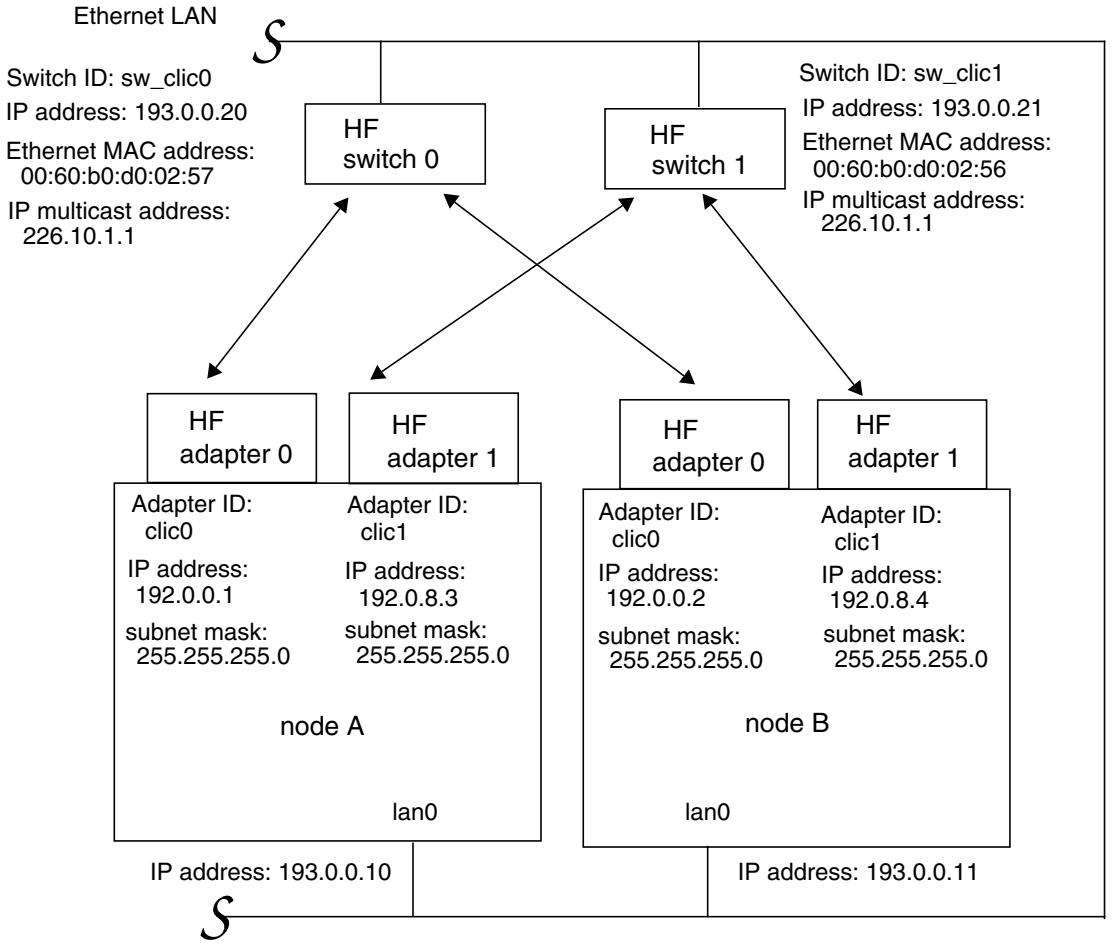

Using the configuration information in Figure 4-1, the information you would specify when you run `clic_init` or SAM on each of the nodes is listed below. This example is not an exact depiction of the prompts produced by `clic_init` nor the fields in SAM, but merely an example of the flow of information input. In addition, you should not try to use the dummy addresses in your actual configuration.

On node A:

1. How many HyperFabric adapters are installed on the node?
2. Do you want this node to interoperate with nodes running any HyperFabric versions earlier than B.11.00.11 or B.11.11.01?
3. What is the IP address of the first adapter (`clic0`)? (`192.0.0.1`)
4. What is the subnet mask of the first adapter? (`255.255.255.0`)

   If you do not specify a value for this, a default mask is chosen. You will most likely just accept the default. However, in this example, we are showing a value for the subnet mask just to illustrate the correlation between the “dummy” information in Figure 4-1 and where that information is specified or generated during `clic_init` and SAM.
5. What is the IP address of the second adapter (`clic1`)? (`192.0.8.3`)
6. What is the subnet mask of the second adapter? (`255.255.225.0`)
7. Do you want to enable switch management? You cannot enable switch management through SAM (you must use the `clic_init` command).

   As stated earlier, HP recommends that you do not enable switch management. However, if you do enable it, you must provide the information in items 8 through 14.
8. If switch management has been enabled, how many switches will be configured? As stated earlier, HP recommends that you do not enable switch management.
9. What is the IP address of the first switch (`sw_clic0`)? (`193.0.0.20`)
10. What is the Ethernet hardware address of the first switch? (`0060b0d00257`)
11. What is the IP address of the second switch (`sw_clic1`)? (`193.0.0.21`)

12. What is the Ethernet hardware address of the second switch? (0060b0d00256)
13. What is the Multicast address for the switches to use? (226.10.1.1)
14. What is the IP address for the LAN card on the same subnet as the switches? (`193.0.0.10`)

    (In Figure 4-1, this is the IP address for `lan0` on node A.)

On node B:

1. How many HyperFabric adapters are installed on the node?
2. Do you want this node to interoperate with nodes running any HyperFabric versions earlier than B.11.00.11 or B.11.11.01?
3. What is the IP address of the first adapter (`clic0`)? (`192.0.0.2`)
4. What is the subnet mask of the first adapter? (`255.255.255.0`)

   If you do not specify a value for this, a default mask is chosen. You will most likely just accept the default. However, in this example, we are showing a value for the subnet mask just to illustrate the correlation between the dummy information in Figure 4-1 and where that information is specified or generated during `clic_init` and SAM.
5. What is the IP address of the second adapter (`clic1`)? (`192.0.8.4`)
6. What is the subnet mask of the second adapter? (`255.255.225.0`)
7. Do you want to enable switch management? You cannot enable switch management through SAM (you must use the `clic_init` command).

   As stated earlier, HP recommends that you do not enable switch management. However, if you do enable it, you must provide the information in items 8 through 14.
8. If switch management has been enabled, how many switches will be configured? As stated earlier, HP recommends that you do not enable switch management.
9. What is the IP address of the first switch (`sw_clic0`)? (`193.0.0.20`)
10. What is the Ethernet hardware address of the first switch? (`0060b0d00257`)
11. What is the IP address of the second switch (`sw_clic1`)? (`193.0.0.21`)

12. What is the Ethernet hardware address of the second switch? (0060b0d00256)
13. What is the Multicast address for the switches to use? (226.10.1.1)
14. What is the IP address for the LAN card on the same subnet as the switches? (`193.0.0.11`)

    (In Figure 4-1, this is the IP address for `lan0` on node B.)

Performing the Configuration

As explained in "Configuration Overview" on page 69, you must create the global configuration file (`/etc/rc.config.d/clic_global_conf`) on each node in the fabric. This consists mostly of specifying HyperFabric adapter-related information. (If you are also going to enable switch management — which HP does not recommend doing — you need to specify additional configuration information.)

NOTE Specifying configuration information adds or changes only the addresses and other information in the global configuration file, based on the information you supply. It does not perform any operations to check the relationships between that information and any physical connections within the fabric.

You need to create the global configuration file in the following situations:

- You have just installed the HyperFabric hardware and software on the system.
- You want to change the information in the HyperFabric global configuration file (see the preceding note).

IMPORTANT Creating the global configuration file also modifies the `/etc/rc.config.d/netconf` file, adding some HyperFabric-related lines that end with the characters `#clic`. These lines are used by the HyperFabric software — and are not comments — so do not remove them from the file.

You can create the global configuration file by using the `clic_init` command (described in "Using the clic_init Command" on page 79) or SAM (described in "Using SAM" on page 81). You cannot enable switch management through SAM (you must use the `clic_init` command).

Using the clic_init Command

Run the `clic_init` command to create the global configuration file.

To view the man page for `clic_init` see “Viewing man Pages” on page 127 of this manual.

To use `clic_init` to configure the Transparent Local Failover feature on HMP, see the section, “Configuring HMP for Transparent Local Failover Support - Using the clic_init command” on page 104.

IMPORTANT

If the global configuration file already exists and you are running `clic_init` again (to change the file), you have the option of retaining or modifying the existing configuration information, in addition to adding new information pertaining to new hardware.

Also, once you have completed your changes and `clic_init` ends its processing, you must stop HyperFabric (by running the `clic_shutdown` command or using SAM) and then start HyperFabric (by running the `clic_start` command or using SAM). Otherwise, your configuration information changes will not take effect. See “Stopping HyperFabric” on page 128 and “Starting HyperFabric” on page 109 for more information.

If you include `/opt/clic/bin` in your `PATH` statement, you can run the command as it is shown below. Otherwise, you must include `/opt/clic/bin` as part of the command name (that is, `/opt/clic/bin/clic_init`).

You must be logged in as `root` to run this command.

The syntax is as follows:

```
clic_init [-c] [-?]
```

where

- `-c` specifies that you want to create the global configuration file. You are prompted for the information described in “Information You Need” on page 71. Note that if the global configuration file already exists (for example, when you are adding an adapter to an existing fabric), `clic_init` prompts you with the existing configuration information. As you are prompted with each piece of information, you can then confirm that you want to keep it or you can change it.
- `-?` displays the online help for `clic_init`.

If you do not specify any of the above parameters, the online help for `clic_init` is displayed.

After you have entered the information for all the adapters in the node and all of the switches (if any) in the fabric, a summary of the configuration information is displayed.

Once `clic_init` has finished, you do one of the following things:

- If you want to configure HyperFabric with ServiceGuard, complete the configuration described in "Configuring HyperFabric with ServiceGuard" on page 87, then run `clic_start` or use SAM to start HyperFabric.
- If you have just created the global configuration file on the local node for the first time (and you are not configuring ServiceGuard), run `clic_start` or use SAM to start HyperFabric.
- If you have just changed an existing configuration file on the node, run `clic_shutdown` or use SAM to stop HyperFabric, and then run `clic_start` or use SAM to start HyperFabric. Until you do those two things, your configuration changes will not take effect.

See "Stopping HyperFabric" on page 128 and "Starting HyperFabric" on page 109 for more information.

Examples of clic_init

Some examples of using the `clic_init` command are shown below.

- **Example 1**

  To create the global configuration file on the local node, issue this command:

  ```
  $ clic_init -c
  ```

- **Example 2**

  To display the online help for the `clic_init` command, issue this command:

  ```
  $ clic_init -?
  ```

  or this command:

  ```
  $ clic_init
  ```

Using SAM

This section describes how to use SAM to configure HyperFabric. For information on how to use SAM to configure and deconfigure local failover feature on HMP, see “Configuring HMP for Transparent Local Failover Support - Using SAM” on page 102 and “Deconfiguring HMP for Local Failover support - Using SAM” on page 103.

IMPORTANT

If the global configuration file already exists, and you are running SAM again (to change the file), you can keep or modify the existing configuration information, in addition to adding new information pertaining to new hardware.

Also, once you’ve completed your changes and SAM ends its processing, you must stop HyperFabric (by running the `clic_shutdown` command or using SAM) and then start HyperFabric (by running the `clic_start` command or using SAM). Otherwise, your configuration information changes will not take effect. See “Stopping HyperFabric” on page 128 and “Starting HyperFabric” on page 109 for more information.

To use SAM to create the global configuration file on an HP 9000 system running HP-UX 11i v2, follow these steps:

Step 1. Start SAM.

Step 2. Select the “Networking and Communications” area.

Step 3. Select “HyperFabric.”

All HyperFabric adapters installed in the system are listed; installed adapters that are not yet configured show `Not Configured` in the “Status” field.

Step 4. Highlight the adapter you want to configure.

Step 5. Pull down the “Actions” menu and select `Configure Adapter`.

Step 6. In the “Configure HyperFabric Adapter” screen, specify information for the following fields:

- `Internet Address`—Required. The IP address of the adapter.

- `Subnet Mask`—Optional. The adapter's subnet mask. If you do not specify this, a default mask is chosen based on the adapter's IP address.
- `Interoperability Enabled`—Required. Whether you want the adapter to be able to interoperate with adapters that are using; any HP-UX 11.0 HyperFabric versions earlier than B.11.00.11 or any HP-UX 11i v1 HyperFabric versions earlier than B.11.11.01. Note that if you select `No`, the HyperFabric software on the system will not be backwards compatible with previous releases. This means you must update all of the other systems in the fabric to the version that is running on the system.

  Default: `No`.

Step 7. Select `OK` (remember, you cannot enable switch management within SAM).

Step 8. Exit SAM.

Once SAM has finished, you do one of the following things:

- If you want to configure HyperFabric with ServiceGuard, complete the configuration described in "Configuring HyperFabric with ServiceGuard" on page 87, then run `clic_start` or use SAM to start HyperFabric.
- If you have just created the global configuration file on the local node for the first time (and you are not configuring ServiceGuard), run `clic_start` or use SAM to start HyperFabric.
- If you have just changed an existing configuration file on the node, run `clic_shutdown` or use SAM to stop HyperFabric, and then run `clic_start` or use SAM to start HyperFabric. Until you do those two things, your configuration changes will not take effect.

See "Stopping HyperFabric" on page 128 and "Starting HyperFabric" on page 109 for more information.

Deconfiguring a HyperFabric Adapter with SAM

To use SAM to deconfigure a HyperFabric adapter on an HP 9000 system running HP-UX 11i v2, complete the following steps:

Step 1. Start SAM.

Step 2. Select the "Networking and Communications" area.

Step 3. Select "HyperFabric."

All HyperFabric adapters installed in the system are listed. Installed adapters that are configured show `Configured` in the "Status" field, and installed adapters that are not yet configured show `Not Configured` in the "Status" field. You can deconfigure only an adapter with a status of `Configured`.

Step 4. Highlight the adapter you want to deconfigure.

Step 5. Pull down the "Actions" menu and select `Deconfigure Adapter`.

Step 6. In the pop-up window, if you want to deconfigure the adapter, select `OK` to confirm it.

If you do not want to deconfigure the adapter, select `Cancel`.

Step 7. If you selected `OK`, the entry for the adapter is deleted from the HyperFabric configuration files (`/etc/rc.config.d/clic_global_conf` and `/etc/rc.config.d/netconf`).

If you selected `Cancel`, you remain in the main "HyperFabric Configuration" screen.

Step 8. Exit SAM.

NOTE If you have configured HMP for Transparent Local Failover support and if you select `Deconfigure Adapter`, HyperFabric will verify if the selected adapter is configured to be part of any card pair. If yes, the user

is informed and the card pair entry is removed from the `/etc/rc.config.d/netconf` and `/etc/rc.config.d/clic_global_conf` files.

Configuring the HyperFabric EMS Monitor

In the HyperFabric version B.11.23.01, the HyperFabric Event Monitoring Service (EMS) monitor allows system administrators to separately monitor each HyperFabric adapter on every node in the fabric, in addition to monitoring the entire HyperFabric subsystem.

The monitor can inform the user if the resource being monitored is UP or DOWN. The administrator defines the condition to trigger a notification (usually a change in interface status). Notification can be accomplished with a SNMP trap or by logging into the syslog file with a choice of severity, or by email to a user defined email address.

To configure the HyperFabric EMS monitor, it is necessary to have the EMS HA monitor product installed (Product Number B7609BA). This product is available on the applications CD media.

Use SAM to initiate monitoring of any particular HyperFabric resource. following the procedure outlined below:

1. Start SAM (Use the online help at any time for details)
2. Select “Resource Management”
3. Select “Event Monitoring Service”
4. Select “Action” and “Add Monitoring Request”
5. Select the location /net/interfaces/clic (class for HyperFabric resources)
6. Select a resource instance (either all instances or a specific instance from the list)
7. Validate your choice by clicking on OK at the bottom of the screen
8. A Monitoring Request Parameters window opens, showing the resource and its status (if All instances have been selected, then no value is displayed)
9. Define a condition that will trigger a notification (for instance, “When Value is”, “equal to”, “UP”)
10. Define a polling interval (default is 300 seconds)
11. Define a way of notification: SNMP trap, log in syslog with a choice of severity, or email to a user defined email address

12. Validate by pressing OK

NOTE Although EMS is able to monitor each HyperFabric adapter on every node in the fabric, as well as the entire HyperFabric subsystem, EMS is not able to monitor HyperFabric switches.

For more detailed information on EMS, including instructions for implementing this feature, see the *EMS Hardware Monitors Users Guide Part Number B6191-90028* September 2001 Edition.

Configuring HyperFabric with ServiceGuard

HyperFabric supports the ServiceGuard HA product.

NOTE If you plan to configure HyperFabric with ServiceGuard, please read this section. Otherwise, skip this section and go on to the next section, “Configuring HMP for Transparent Local Failover Support” on page 96.

ServiceGuard lets you create HA clusters of HP 9000 server systems. Within the cluster, ServiceGuard allows you to group your application services (individual HP-UX processes) into packages. In the event of a single service, node, network, or other resource failure, ServiceGuard can transfer control of the package to another node in the cluster, allowing services to remain available with minimal interruption.

ServiceGuard directly monitors cluster nodes, LAN interfaces, and services, which are the individual processes within an application. In addition, specialized monitors might be supplied by the developers of other components. The HyperFabric monitor is supplied with the HyperFabric product and is installed with it. To use the HyperFabric monitor with ServiceGuard, you configure the monitor as an ServiceGuard package dependency.

Although HyperFabric can be used by an application within a package to communicate with other nodes, it is not possible to use HyperFabric as a heartbeat LAN. So, in a package control script, do not specify HyperFabric IPs/subnets in the lines that contain the keywords `IP[`*n*`]` and `SUBNET[`*n*`]`. Also, `cmquerycl` will not “discover” and report HyperFabric IPs and subnets.

After you have configured HyperFabric as a package dependency, ServiceGuard’s package manager calls the **Event Monitoring Service** (**EMS**) to launch an external monitor for HyperFabric. The package will not start unless the monitor reports that HyperFabric is available, and the package will fail when HyperFabric’s status is DOWN (that is, when all HyperFabric adapters on a node become non-functional).

Complete instructions for configuring ServiceGuard clusters and packages are provided in the manual, “*Managing MC/ServiceGuard*”.

Figure 4-2 below shows a HyperFabric switch configuration with ServiceGuard. This example shows a four-node configuration with two HyperFabric switches, and redundant heartbeat Ethernet LANs.

NOTE Because the HyperFabric network does not currently support ServiceGuard heartbeat connections, you must use an alternative type of connection for the heartbeat, such as FDDI, Token Ring, 100BaseT, or Ethernet (as shown in Figure 4-2).

Figure 4-2 An ServiceGuard Configuration (with Two HyperFabric Switches)

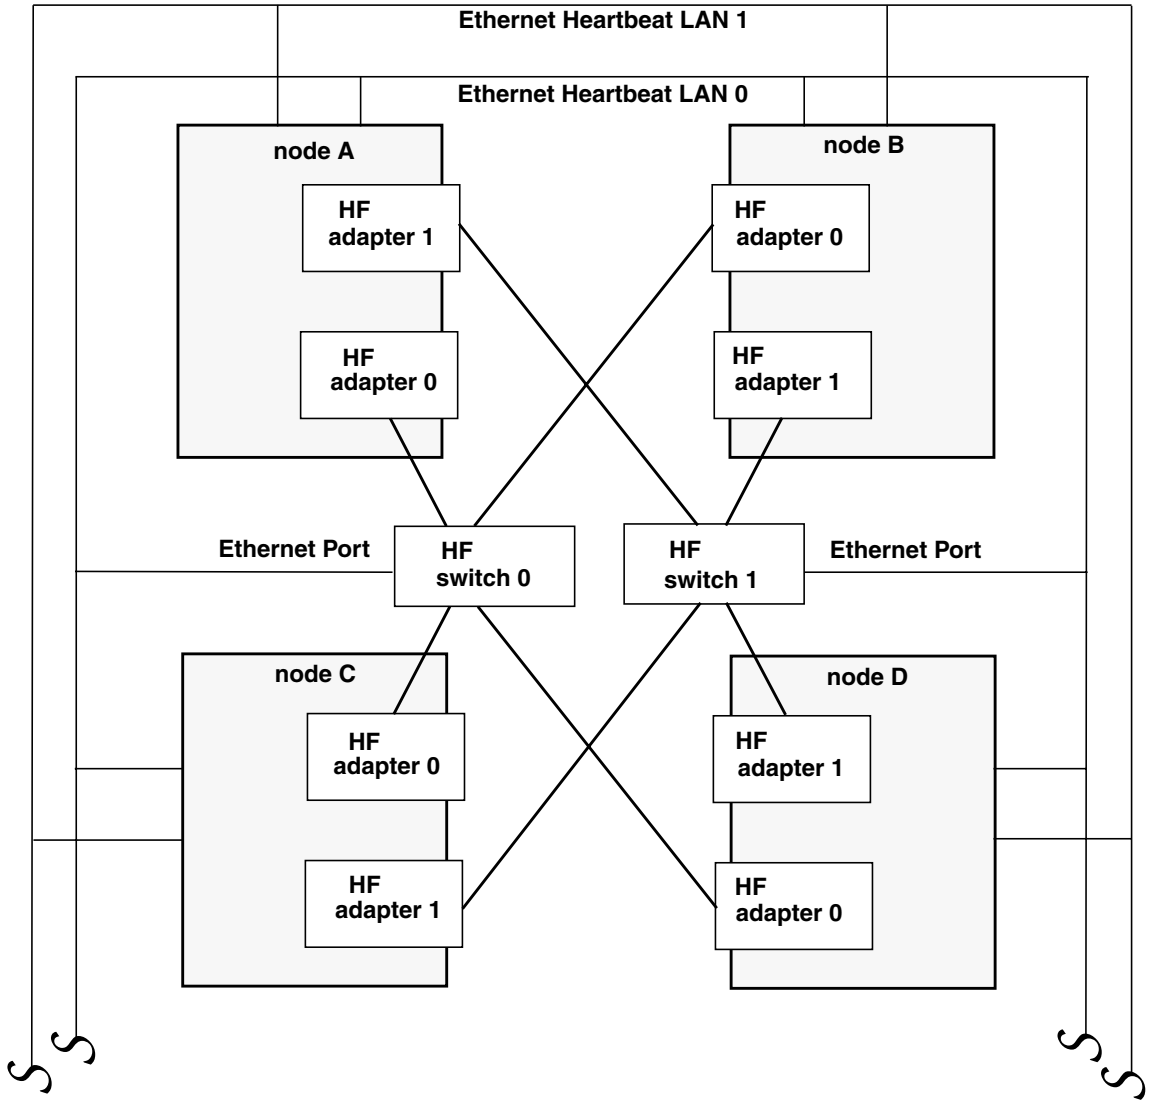

How HyperFabric Handles Adapter Failures

HyperFabric adapters are handled differently than other types of networking adapters (such as Ethernet, FDDI, and Fibre Channel) in the ServiceGuard environment. In the non-HyperFabric cases, two network links are in a node, and one will be active and one will be idle or in standby. In the case of an active link failure, ServiceGuard is notified and the network traffic is switched to the standby adapter (which then becomes active).

However, in the case of HyperFabric, if two adapters are in a node, both will be active. If one active HyperFabric adapter fails, its network traffic is switched to the other active HyperFabric adapter in the node. (Throughput might be slower because only one active adapter is now handling the network traffic.) This rearrangement is handled by the HyperFabric software, and ServiceGuard is not notified. However, note that if all of the HyperFabric adapters fail, HyperFabric does notify ServiceGuard. In both cases, though, the events are logged to `/var/adm/clic_log` and `/var/adm/syslog.log`.

Example 1:

This example, illustrated by Figure 4-3 below, presents an HA configuration using ServiceGuard with HyperFabric. Both of the HyperFabric adapters are active on node A. The HyperFabric Resource Monitor reports the active status of the HyperFabric resource to the Event Monitoring Service (EMS), which lets ServiceGuard know that the HyperFabric resource is available to Packages A and B.

Figure 4-3 Node with Two Active HyperFabric Adapters

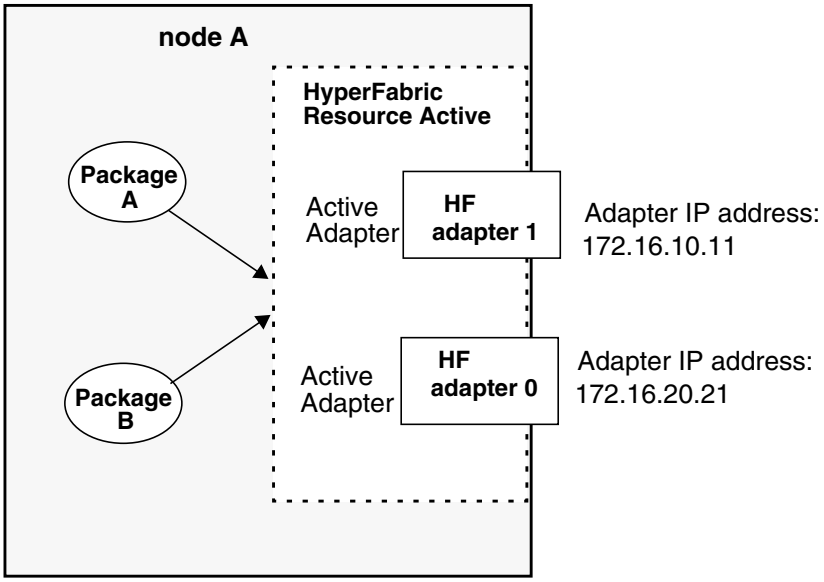

Example 2:

This example, illustrated by Figure 4-4 below, shows the same node after the failure of one of the HyperFabric adapters. The remaining adapter in node A is now handling all HyperFabric network traffic for the node. Because the HyperFabric resource is still available, ServiceGuard has not been notified; HyperFabric handles the local HyperFabric adapter failover. However, the failure of adapter 1 has been logged to `/var/adm/clic_log`.

Figure 4-4 Node with One Failed HyperFabric Adapter

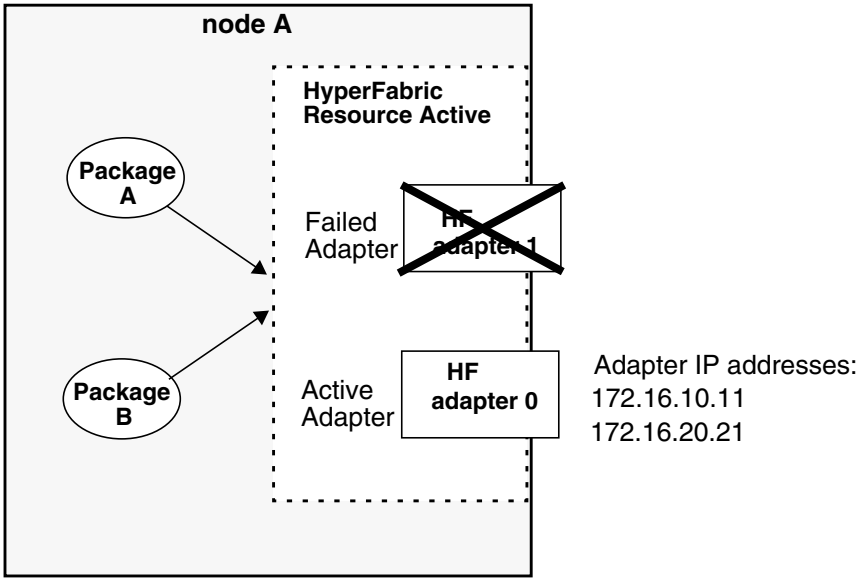

After the failover, if you issue a `netstat -in` command, you will see that an IP address is still assigned to each adapter. For example:

```
Name    MTU      network      Address        Ipkts  Opkts
clic1   31744    172.16.10.0  172.16.10.11   711    12
clic0   31744    172.16.20.0  172.16.20.21   1222   333
```

Example 3:

This final example, illustrated by Figure 4-5 below, shows a situation in which all of the HyperFabric adapters on node A fail. The HyperFabric Resource Monitor reports to the Event Monitoring Service (EMS). The EMS then notifies the ServiceGuard `cmcld` daemon that the HyperFabric resource on node A is unavailable. Because HyperFabric is configured as a package dependency for Packages A and B, ServiceGuard causes the packages to failover to node B. In a four-node configuration (note that only two nodes are shown in Figure 4-5 below), Packages A and B can continue to communicate through the HyperFabric network with the other active nodes in the ServiceGuard cluster.

Figure 4-5 When All HyperFabric Adapters Fail

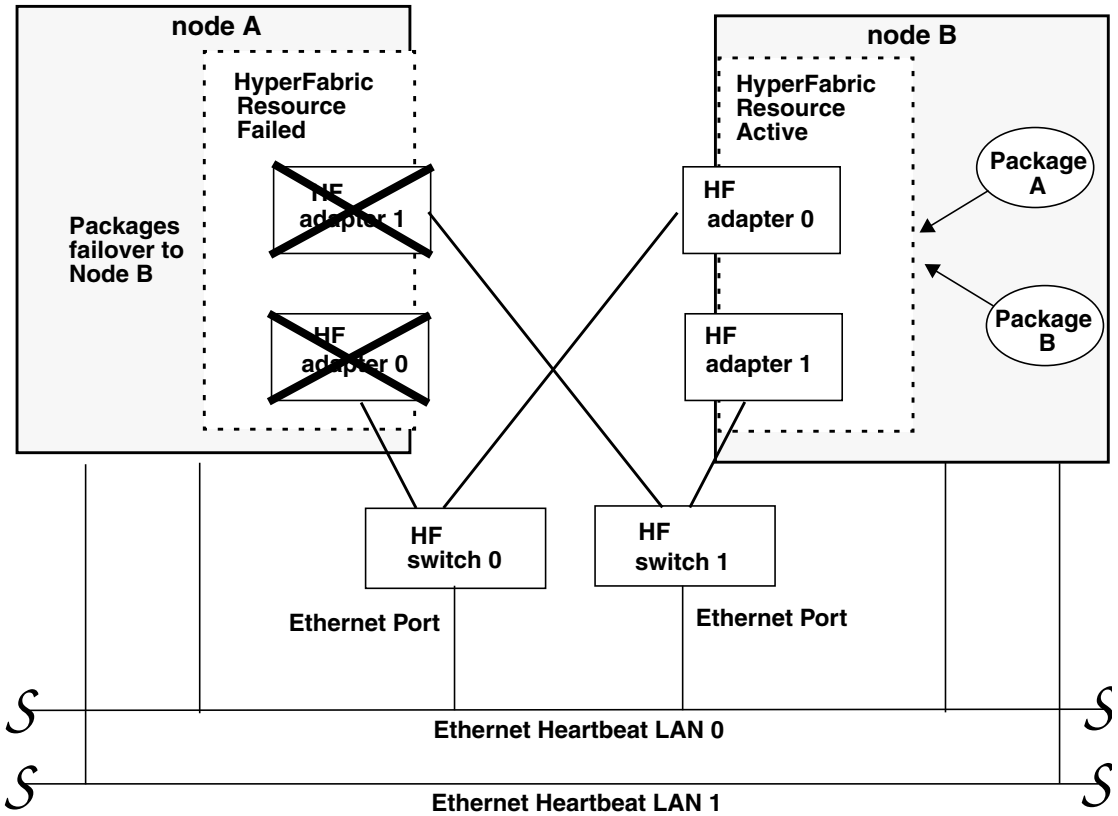

Configuring HyperFabric with the ServiceGuard Resource Monitor

You can configure the HyperFabric Resource Monitor with ServiceGuard in either of these ways:

- Editing an ASCII file.
- Using the SAM GUI.

For more details, please see the manual *Using EMS HA Monitors*.

NOTE You should configure HyperFabric with ServiceGuard before running the `clic_start` command or using SAM to start HyperFabric.

Configuring ServiceGuard with HyperFabric Using the ASCII File

When using the ServiceGuard commands (for example, `cmapplyconf`) to specify the use of the HyperFabric Resource Monitor, the section of the package ASCII configuration file that has the keyword `RESOURCE_NAME` must be uncommented and set to the following values:

```
RESOURCE_NAME                /net/interfaces/clic/status

RESOURCE_POLLING_INTERVAL    10

RESOURCE_UP_VALUE            =UP
```

Configuring ServiceGuard with HyperFabric Using SAM

You must perform the following steps when using SAM to configure the HyperFabric Resource Monitor with ServiceGuard:

```
sam
 Clusters
   High Availability Clusters
     Cluster Configuration (go through all the steps to create
a
     cluster)
     Package Configuration
       Create/Add Package (if creating new packages)
```

```
          Specify Package Name and Nodes
          Specify Package SUBNET Address
          Specify Package Services
          Specify Package Failover Options
          Specify Package Control Script Location
          Specify Package Control Script Information

  Specify Package Resources Dependencies
 Add
       Resource Name
        (Navigate the Resource Subclass by double-clicking on
         /net until /net/interfaces/clic/ status shows up in
the
         selection box Resource Name,then select it and click
OK.)

Resource Parameters

- Input the Resource Polling Interval (for example, 10
seconds).

- Select "UP" from the "Available Resource Values" and click
"Add".

- Click OK to accept the values.
```

Configuring ServiceGuard for HyperFabric Relocatable IP Addresses

If you are using HyperFabric version B.11.00.05, B.11.11.00, or later, and you want to use relocatable IP addresses, configure the relocatable IP addresses with the `IP[n]` command in the package control script.

For example, to configure the relocatable address `192.0.0.3` for adapter 0 and `192.0.8.5` for adapter 1, specify this:

```
IP[0]= 192.0.0.3

IP[1]= 192.0.8.5
```

Configuring HMP for Transparent Local Failover Support

HMP supports Local Failover in the HyperFabric version B.11.23.01.

If a HyperFabric resource (adapter, cable, switch or switch port) fails in a cluster, HMP transparently fails over traffic (Local Failover) using another available resource from the card pair. A card pair can be defined as a logical entity comprising of a pair of HF2 adapters on a HP 9000 node. For example, if there are four HF2 adapters installed and configured in a node, then there would be two card pairs.

IMPORTANT

Remember the following points while configuring HMP for Local Failover support:

- Only Oracle applications can make use of the local failover feature. Other middleware like MPI can continue using HMP without local failover support.
- HMP supports the local failover configuration in HyperFabric version B.11.23.01.
- All nodes in the cluster need to be configured either in the local failover mode or the non-local failover mode. While using `clic_init`, if you answer 'y' to the question, "Do you want to configure Local Failover on this node?", then you have configured HMP for the local failover mode. Otherwise, your HMP configuration is for the non-local failover mode. Do not mix these two modes. For any incorrect configurations, HP recommends that you `clic_shutdown` and `clic_start` the cluster.
- The Transparent Local Failover feature over HMP is supported only in a switch-based environment. If two nodes are connected through multiple point-to-point links, local failover cannot be achieved.
- There must be even number of HF2 adapters on any given node, and all adapters installed must be configured. In addition, all the adapters must belong to a card pair.
- HMP can fail over traffic only between adapters that belong to the same card pair.

- HMP does not support backward compatibility in the local failover and non-local failover mode. However, TCP/UDP/IP supports backward compatibility and interoperability.
- When HMP is configured in the local failover mode, all the resources in the cluster are utilized. If a resource fails in the cluster and is restored, HMP does not utilize that resource until another resource fails.
- Before running `clic_start` on all the nodes in the cluster, ensure that all the configured cards are connected in the cluster. In other words, before running the `clic_start` command, verify that all the cables are connected to the adapters and switches.
- If a resource in the cluster fails and is restored, perform the following steps to ensure full utilization of that resource:
  - Shutdown Oracle RAC
  - Execute the `clic_shutdown` command
  - Execute the `clic_start` command
  - Restart Oracle RAC
- After executing the clic_start command on all node in the cluster, ensure that you run Oracle RAC only after one minute (approximately).
- If all the trunks between the switches are down, then execute `clic_shutdown` followed by `clic_start` on all the nodes in the cluster, after replacing at least one trunk between the switches.
- Maintain at least one trunk between the switches when Oracle RAC is running in the cluster.
- If any of the nodes is shut down for administrative reasons, shut down Oracle RAC before executing `clic_shutdown` on that node. In such a case, Oracle RAC continues to be operational on all the other nodes in the cluster. When you bring up the node that was shut down, it joins the cluster.

How Transparent Local Failover Works

Consider a hypothetical HyperFabric configuration in a 4-node cluster, with each node having two adapters (see Figure 4-6). In this configuration, there is no single point of failure, and all adapters that are installed on any given node are configured as part of a card pair.

Figure 4-6 A Configuration supporting Local Failover

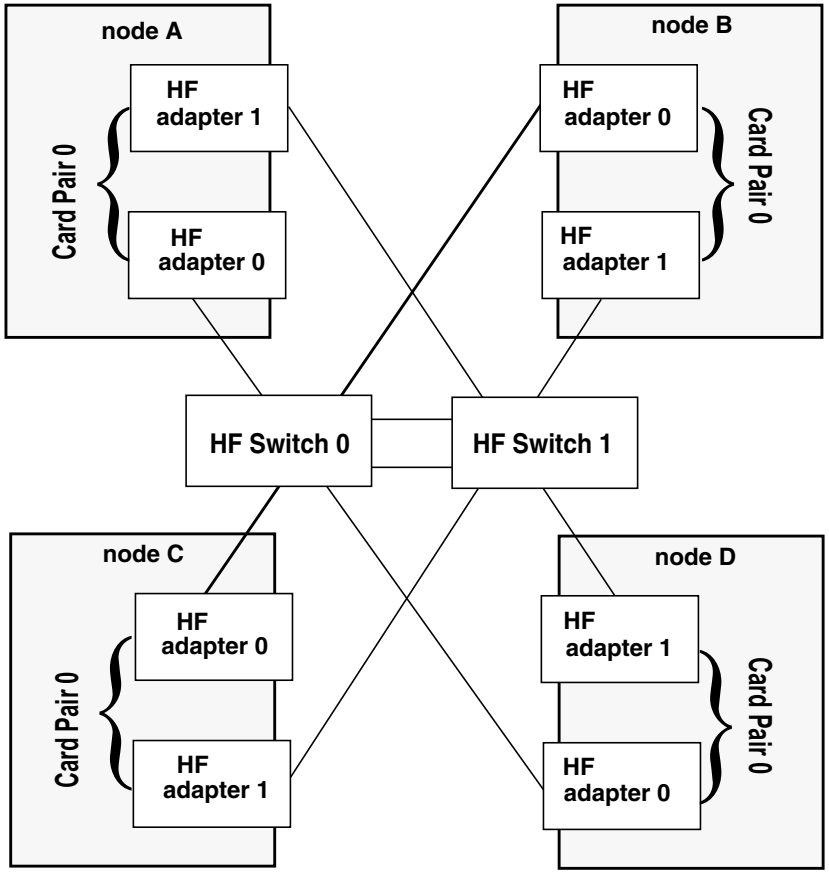

The details of how HyperFabric handles failures of the adapters, links, switch ports, switches, and cables between switches in a cluster are discussed in the following sections.

Case 1: Adapter, Link or Switch Port Failure (see Figure 4-7)

If an adapter or a link or a switch port fails, HMP transparently fails over traffic through the other available link.

Figure 4-7 Adapter, Link or Switch Port Failover

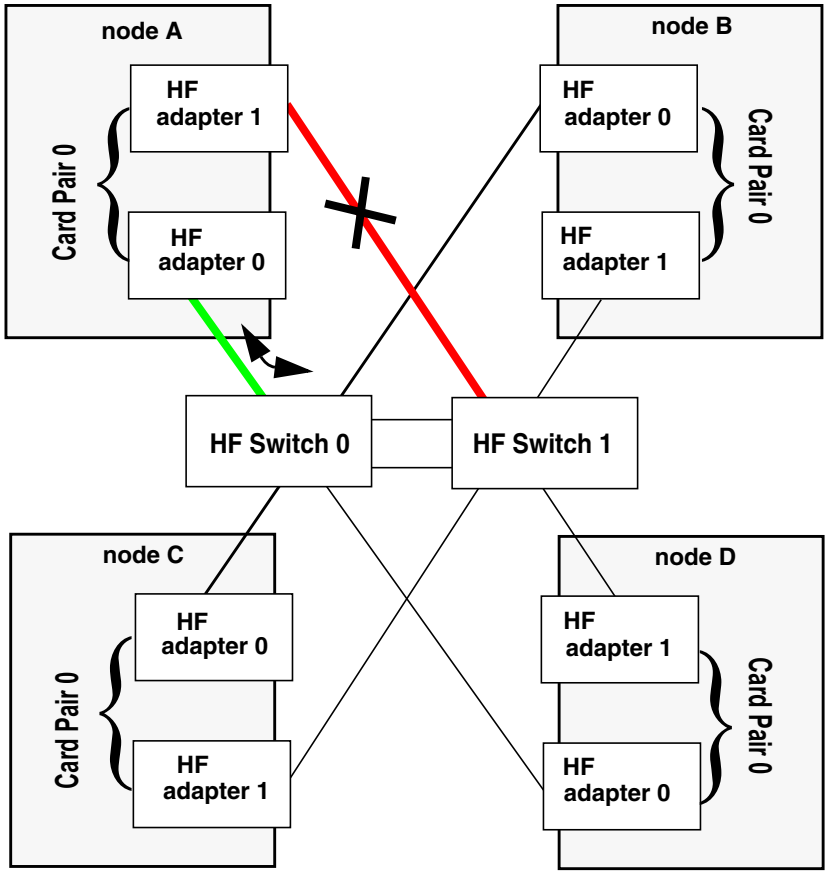

Case 2: Switch Failure (see Figure 4-8)

Consider the following illustration where node A is connected to node D with traffic being routed through the HF adapter 1 on both the nodes (A and D), and the HF switch 1 fails. HMP transparently fails over traffic through the other available switch (HF switch 0). This is possible only if at least one adapter of the card pairs on both the nodes (HF adapter 0 of the Card Pair 0 on nodes A and D) is physically reachable (through the HF switch 0).

Figure 4-8 Switch Failover

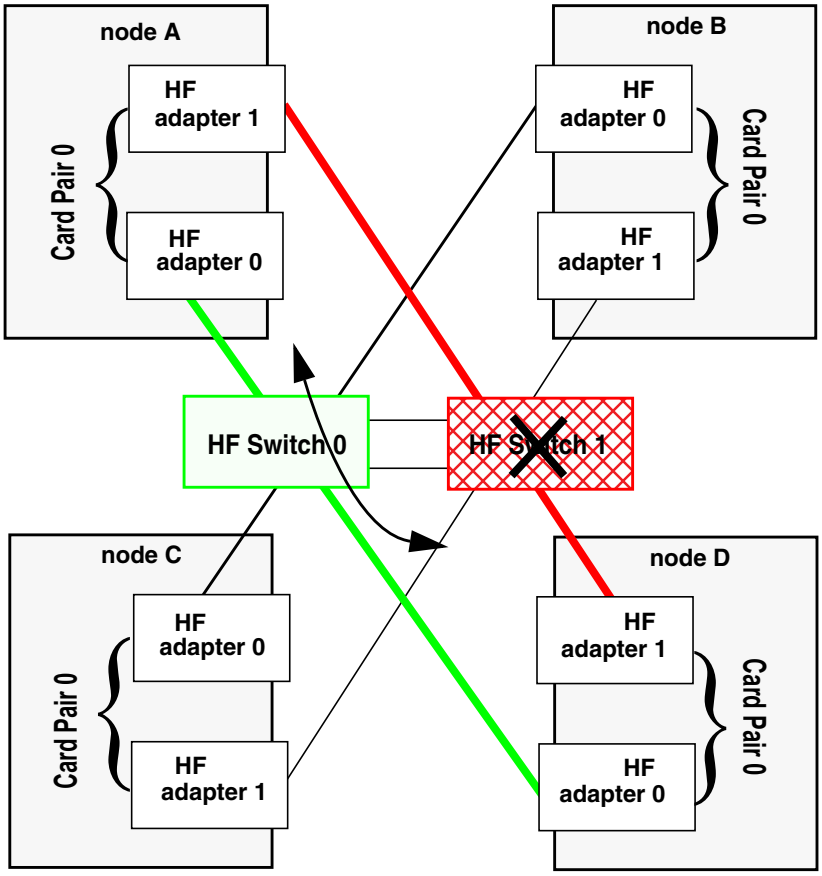

Thus, if a switch fails, HMP transparently fails over traffic only if at least one member of the card pair is physically reachable through the other switch.

Case 3: Cable Failure Between Two Switches (see Figure 4-9)

If a cable between two switches fails, HMP traffic fails over to the other available cable between those two switches.

Figure 4-9 Cable Failover Between Two Switches

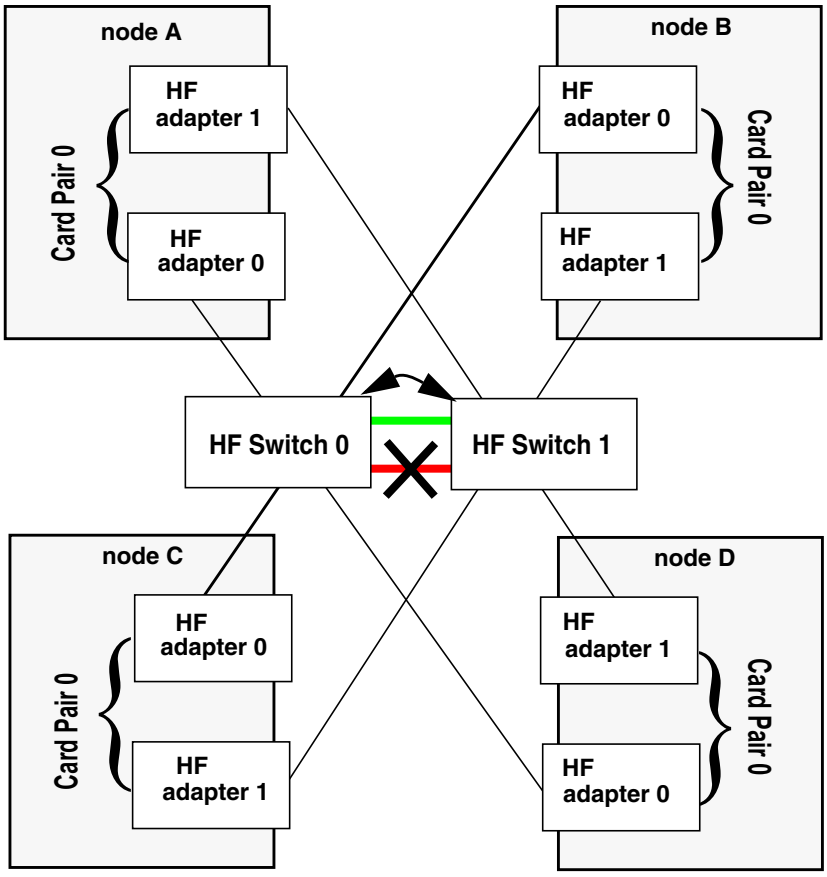

Configuring HMP for Transparent Local Failover Support - Using SAM

To use SAM to configure HMP for Local Failover Support, complete the following steps:

Step 1. Start SAM.

Step 2. Select the “Networking and Communications” area.

Step 3. Select “HyperFabric.”

All HyperFabric adapters installed in the system are listed; installed adapters that are not yet configured show `Not Configured` in the “Status” field. Before proceeding to the next step, configure all the available adapters.

Step 4. Pull down the "Actions" menu and select "Configure Local Failover for HMP". SAM now verifies the following:

- If not all the adapters are configured
- If the number of adapters configured are odd
- If “Interoperability Enabled” is set to YES

If any of the above conditions is true, then SAM displays an appropriate error message. Otherwise, the “Configure Local Failover for HMP” window pops up.

Step 5. In this window, specify the configured adapters for each card pair and press `OK`.

On pressing `OK`, SAM verifies the following:

- If HyperFabric subsystem is not running on the machine
- If all card pairs are not configured
- If you have chosen the same adapter for different card pairs

If any of the above conditions is true, SAM displays an appropriate error message. Otherwise, the `/etc/rc.config.d/clic_global_conf` file is updated with information about the configured card pairs. If Card-Pair 0 comprises

of adapters clic0 and clic1 and if the Card-Pair 1, comprises of adapters clic2 and clic3, then the following entries are added to the `clic_global_conf` file.

```
CARD_PAIR[0] = clic0-clic1
CARD_PAIR[1] = clic2-clic3
```

If you press `Cancel`, you remain in the main "HyperFabric Configuration" screen.

If you press `Help`, help text for this task appears.

Step 6. Exit SAM.

NOTE To view the card pair information from the `/etc/rc.config.d/clic_global_conf file`, select "View Local Failover for HMP..." option in the "HyperFabric Configuration" screen.

Deconfiguring HMP for Local Failover support - Using SAM

To use SAM to deconfigure HMP for Local Failover support, complete the following steps:

Step 1. Start SAM.

Step 2. Select the "Networking and Communications" area.

Step 3. Select "HyperFabric".

All HyperFabric adapter card pairs installed in the system are listed.

Step 4. Pull down the "Actions" menu and select `Deconfigure Local Failover for HMP` to remove all the card pair entries from the `/etc/rc.config.d/clic_global_conf` file. In the next confirmation window that pops-up, select `YES` to confirm it.

If you select `NO`. you remain in the main "HyperFabric Configuration" screen.

Step 5. Exit SAM.

Configuring HMP for Transparent Local Failover Support - Using the `clic_init` command

You can configure the Transparent Local Failover feature of HMP using clic_init also.

Let us consider the following example where we have discussed the configuration in detail. This example uses some “dummy” (that is, not valid) addresses to the components in Figure 4-10. The dummy addresses are used only to show the flow of the information provided as input to the `clic_init` command and SAM. Do not try to use these addresses in your configuration.

Figure 4-10 Configuring the Transparent Local Failover feature

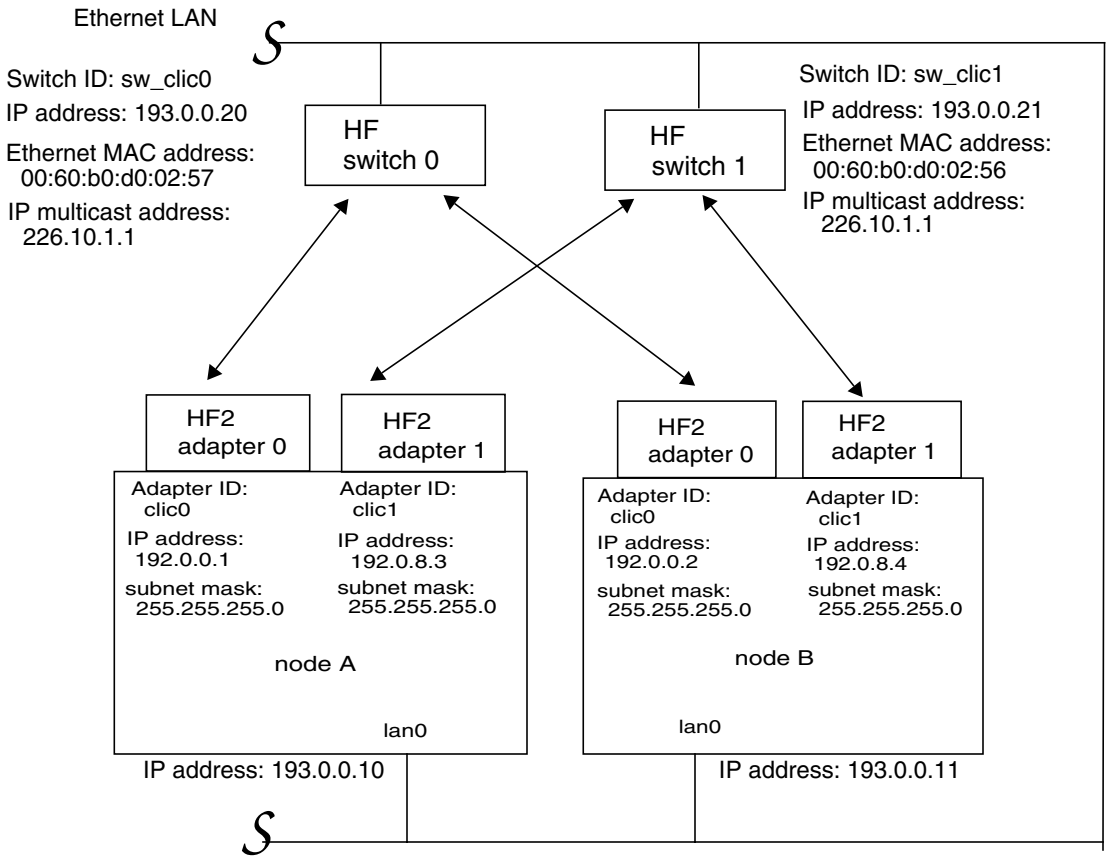

Using the configuration information in Figure 4-10, the information you would specify when you run `clic_init` on each of the nodes is listed below. This example is not an exact depiction of the prompts produced by `clic_init`, but merely an example of the flow of information input. In addition, you should not try to use the dummy addresses in your actual configuration.

On node A:

1. How many HyperFabric adapters are installed on the node?
2. Do you want this node to interoperate with nodes running any HyperFabric versions earlier than B.11.00.11 or B.11.11.01? (n)

   You must answer 'no' if you want to run applications using HMP (Local Failover or Non-Local Failover) for communication over HyperFabric. In that case, all nodes in the cluster must be running version B.11.00.11 (or) B.11.11.01 (or) later version of HyperFabric software.
3. What is the IP address of the first adapter (`clic0`)? (`192.0.0.1`)
4. What is the subnet mask of the first adapter? (`255.255.255.0`)

   If you do not specify a value for this, a default mask is chosen. You will most likely just accept the default. However, in this example, we are showing a value for the subnet mask just to illustrate the correlation between the "dummy" information in Figure 4-10 and where that information is specified or generated during `clic_init` and SAM.
5. What is the IP address of the second adapter (`clic1`)? (`192.0.8.3`)
6. What is the subnet mask of the second adapter? (`255.255.255.0`)
7. Do you want to configure Local Failover on this node? (y)

   Enter 'y' if you are using Oracle RAC with HMP. For MPI applications, enter 'n'.
8. Select any two of the following clic adapters for CARD_PAIR[0]:
   clic0
   clic1

   Enter the first clic adapter from above listed adapters: (clic0)

   Enter the second clic adapter from above listed adapters: (clic1)
9. Do you want to enable switch management? (n)

On node B:

1. How many HyperFabric adapters are installed on the node?
2. Do you want this node to interoperate with nodes running any HyperFabric versions earlier than B.11.00.11 or B.11.11.01? (n)

   You must answer 'no' if you want to run applications using HMP (Local Failover or Non-Local Failover) for communication over HyperFabric. In that case, all nodes in the cluster must be running version B.11.00.11 (or) B.11.11.01 (or) later version of HyperFabric software.
3. What is the IP address of the first adapter (`clic0`)? (`192.0.0.2`)
4. What is the subnet mask of the first adapter? (`255.255.255.0`)

   If you do not specify a value for this, a default mask is chosen. You will most likely just accept the default. However, in this example, we are showing a value for the subnet mask just to illustrate the correlation between the dummy information in Figure 4-10 and where that information is specified or generated during `clic_init` and SAM.
5. What is the IP address of the second adapter (`clic1`)? (`192.0.8.4`)
6. What is the subnet mask of the second adapter? (`255.255.255.0`)
7. Do you want to configure Local Failover on this node? (y)

   Enter 'y' if you are using Oracle RAC with HMP. For MPI applications, enter 'n'.
8. Select any two of the following clic adapters for CARD_PAIR[0]:
   clic0
   clic1

   Enter the first clic adapter from above listed adapters: (clic0)

   Enter the second clic adapter from above listed adapters: (clic1)
9. Do you want to enable switch management? (n)

5 Managing HyperFabric

This chapter contains the following sections that give information about managing HyperFabric:

Starting HyperFabric

HyperFabric is started in one of these three ways:

- As part of the normal local node boot process (HP 9000 system).
- By running the HyperFabric `clic_start` command (described below).
- By starting HyperFabric through SAM (described in “Using SAM” on page 110).

HyperFabric needs to be started in the following situations:

- If HyperFabric hardware and software have just been installed on the system and the `clic_init` command or SAM has been used to configure the HyperFabric adapters on this node.
- If the HyperFabric configuration has been changed by using the `clic_init` command or using SAM. In this situation, you must have run `clic_shutdown` or used SAM to stop HyperFabric, before restarting HyperFabric.
- If a new HyperFabric adapter has been added to a system online and configured using `clic_init`. In this situation, it is not necessary to run `clic_shutdown` before running `clic_start` (see “Online Addition and Replacement” on page 44).

NOTE

Starting HyperFabric launches the HyperFabric CLuster InterConnect (CLIC) daemon (`clic_mgmtd`). This daemon process must be running for the HyperFabric product to operate correctly. It is possible that other daemons will be running, but it is essential that at least one CLIC daemon is running. To check if a CLIC daemon is running, use the following command:

```
ps -ef | grep clic
```

If the CLIC daemon is not running, start the HyperFabric subsystem by executing the following command:

```
/opt/clic/bin/clic_start
```

Using the clic_start Command

Run the `clic_start` command on each node to start the HyperFabric management process on that node.

If you include `/opt/clic/bin` in your `PATH` statement, you can run the command as it is shown below. Otherwise, you must include `/opt/clic/bin` as part of the command name (that is, `/opt/clic/bin/clic_start`).

You must be logged in as `root` to run this command.

The syntax is as follows:

```
clic_start
```

The `clic_start -?` command can be issued to display the online help for `clic_start`, or look at the `clic_start (1m)` man page by issuing the `man clic_start` command.

If HyperFabric is already running, you will receive an informational (FYI) message telling you so. Your reaction to this message depends on the situation:

- If you have simply forgotten (or did not know) that HyperFabric was already running, you do not have to do anything.
- If you have changed the HyperFabric configuration with the `clic_init` command or SAM, you must stop HyperFabric (by running the `clic_shutdown` command or using SAM) and then start HyperFabric (by running the `clic_start` command or using SAM). See either “Using the clic_shutdown Command” on page 128 or “Using SAM” on page 129, whichever is appropriate.

Using SAM

To use SAM to start HyperFabric on an HP 9000 system running HP-UX 11i v2, complete the following steps:

Step 1. Start SAM.

Step 2. Select the “Networking and Communications” area.

Step 3. Select “HyperFabric”.

Step 4. Pull down the “Actions” menu and select `Start HyperFabric`.

When HyperFabric starts, a confirmation message displays. Also, the status “HyperFabric: Running” is displayed above the adapter configuration area of the screen.

Step 5. Exit SAM.

Verifying Communications within the Fabric

You can verify the communications within the fabric by running the `clic_probe` command, which is described below. You can also use `clic_probe` to verify the status of specific adapters.

IMPORTANT

You should also check your `/etc/hosts` file—when you are using files for host name look up—to ensure that the entries for all of the systems are in the correct format: the official host name, which is the full domain extended host name, and any alias names. For example:

```
IP_address   bently6.corp3.com   bently6
IP_address   bently4.corp7.com   test1
IP_address   bently2.corp4.com   test3
```

The clic_probe Command

Run the `clic_probe` command to send 256-byte packets to verify the link out to and back from a specific destination, optionally using a specific adapter for the verification. The destination can be either a node or a switch (if a switch is part of the fabric).

If you include `/opt/clic/bin` in your `PATH` statement, you can run the command as it is shown below. Otherwise, you must include `/opt/clic/bin` as part of the command name (that is, `/opt/clic/bin/clic_probe`).

You do not have to be logged in as `root` to run this command.

The syntax is as follows:

```
clic_probe node_name [-c adapter_ID]
           [-c adapter_ID -r VRID switch_hopcount]
           [-l -c adapter_ID] [-s -c adapter_ID]
           [-p packet_count] [-?]
```

Note that some of the lines in the above syntax are indented for readability purposes only. When you actually type the command, you do not indent anything.

The command parameters are as follows:

- *node_name* specifies the node you want to verify. This value is conditionally required—you must specify it when you are verifying traffic to a remote node, unless you use the -r parameter (described below).
- -c specifies that you want to use the adapter identified by *adapter_ID* for the verification.
- -r specifies that *VRID switch_hopcount* is the routing information for the adapter. To determine the *VRID* and *switch_hopcount* to specify, first run the clic_stat -d VRID command (see "The clic_stat Command" on page 117). Note that if you specify this parameter (-r *VRID switch_hopcount*), you must also specify the -c *adapter_ID* parameter (described above).
- -l specifies that you want to do local loopback testing on a particular adapter. Note that if you specify this parameter (-l), you must also specify the -c *adapter_ID* parameter (described above).
- -s specifies that you want to loopback at the switch port attached to a particular adapter. Note that if you specify this parameter (-s), you must also specify the -c *adapter_ID* parameter (described above).
- -p specifies that you want to send *packet_count* number of 256-byte packets. *packet_count* can be any positive integer. This parameter is useful for building scripts for debugging or for hardware verification. If you do not specify this parameter, one packet is sent each second, until you stop the command with a **CTRL-C**.
- -? displays the online help for clic_probe.

If you do not specify any of the above parameters, the online help for clic_probe is displayed.

NOTE

Also see the `clic_diag` command to:

Probe a specific remote node.

Dump and format trace data.

Set the tracing level for the HyperFabric software and firmware.

The `clic_diag` command is detailed in the section, "Running Diagnostics" on page 133.

Examples of clic_probe

Some examples of using `clic_probe` are shown below.

- **Example 1**

  If the local node is `bently6` and you want to send five packets to verify that the adapter `clic0` (which is on `bently6`) is able to handle traffic, issue this command:

  `clic_probe -l -c clic0 -p 5`

  The generated output could look like this:

```
CLIC_PROBE: 256 byte packets
Local Loopback: Source and Target Adapter ID: bently6.corp
3.com:clic0

256 bytes:     seq_num = 1.    Packet Acknowledged.
256 bytes:     seq_num = 2.    Packet Acknowledged.
256 bytes:     seq_num = 3.    Packet Acknowledged.
256 bytes:     seq_num = 4.    Packet Acknowledged.
256 bytes:     seq_num = 5.    Packet Acknowledged.
--------- bently6.corp3.com CLIC_PROBE Statistics --------
5 packets transmitted, 5 packets received, 0% packet loss.
```

- **Example 2**

  If the local node is `bently6`, and you want to verify communications with the remote node `bently4`, issue this command:

  `clic_probe bently4`

```
CLIC_PROBE: 256 byte packets
  Source adapter id: bently6.corp3.com:clic0
  Target adapter id: bently4.corp7.com:clic3

256 bytes:     seq_num = 1.    Packet Acknowledged.
256 bytes:     seq_num = 2.    Packet Acknowledged.
256 bytes:     seq_num = 3.    Packet Acknowledged.
256 bytes:     seq_num = 4.    Packet Acknowledged.
256 bytes:     seq_num = 5.    Packet Acknowledged.
256 bytes:     seq_num = 6.    Packet Acknowledged.
256 bytes:     seq_num = 7.    Packet Acknowledged.
256 bytes:     seq_num = 8.    Packet Acknowledged.
--------- bently6.corp3.com CLIC_PROBE Statistics --------
8 packets transmitted, 8 packets received, 0% packet loss.
```

- **Example 3**

  If the local node is `bently6`, and you want to send five packets to verify communications with the remote node `bently7`, using the adapter `clic0` (which is on `bently6`), issue this command:

  `clic_probe bently7 -c clic0 -p 5`

```
CLIC_PROBE: 256 byte packets
  Source adapter id: bently6.corp3.com:clic0
  Target adapter id: bently7.corp4.com:clic1

256 bytes:     seq_num = 1.    Packet Acknowledged.
256 bytes:     seq_num = 2.    Packet Acknowledged.
256 bytes:     seq_num = 3.    Packet Acknowledged.
256 bytes:     seq_num = 4.    Packet Acknowledged.
256 bytes:     seq_num = 5.    Packet Acknowledged.
--------- bently7.corp4.com CLIC_PROBE Statistics --------
5 packets transmitted, 5 packets received, 0% packet loss.
```

- **Example 4**

  If the local node is `bently6`, and you want to send five packets to verify communications with the remote node `bently7`, using the adapter `clic0` (which is on `bently6`) and the route identified by VRID `194` and switch hopcount `1`, issue this command:

  `clic_probe -c clic0 -r 194 1 -p 5`

  (Remember, because you specified the `-r` *`VRID switch_hopcount`* parameter, you do not need to also specify the *`node_name`*.)

The generated output could look like this:

```
CLIC_PROBE: 256 byte packets sent
  Source adapter id: bently6.corp3.com:clic0
  Target adapter id: bently7.corp4.com:clic1

256 bytes:     seq_num = 1.    Packet Acknowledged.
256 bytes:     seq_num = 2.    Packet Acknowledged.
256 bytes:     seq_num = 3.    Packet Acknowledged.
256 bytes:     seq_num = 4.    Packet Acknowledged.
256 bytes:     seq_num = 5.    Packet Acknowledged.
--------- bently7.corp4.com CLIC_PROBE Statistics --------
5 packets transmitted, 5 packets received, 0% packet loss.
```

Note that the VRID you specified (`194`) actually went to the adapter `clic1` on `bently7`. And, as explained earlier, you run the `clic_stat -d VRID` command to determine the VRID and switch hopcount to specify.

Displaying Status and Statistics

You can get the status of and statistics associated with many of the HyperFabric components by using the `clic_stat` command, which is described below.

The clic_stat Command

The following list contains some of the information that the `clic_stat` command provides:

- The current fabric map, in textual format.
- The status of one or more HyperFabric adapters.
- The global configuration information for each HyperFabric adapter and switch (if the fabric contains switches). The information includes the firmware type (8-bit or 32-bit), which is used for interoperability purposes.
- The card pair information, if you configure HyperFabric in the local failover mode (applicable to applications using HMP).

`The clic_stat` command can also be used to enable or disable performance statistics gathering for the DLPI driver, the firmware and HMP.

All of the statistics that can be displayed using the `clic_stat` command are documented in the `clic_stat (1M)` man page on HP-UX 11.0 and on later HP-UX releases.

If you include `/opt/clic/bin` in your `PATH` statement, you can run the command as it is shown below. Otherwise, you must include `/opt/clic/bin` as part of the command name (that is, `/opt/clic/bin/clic_stat`).

To use some of this command's parameters, you must be logged in as `root` (see each parameter's description below).

The syntax is as follows:

```
clic_stat [-p perf_level] [-d display_level] [-c adapter_ID]
          [-n nodename] [-?]
```

Note that the second line in the above syntax is indented for readability purposes only. When the command is typed there should not be any indentation.

The command parameters are as follows:

- `-p` enables/disables performance statistics gathering according to the value of *`perf_level`*, which is one of the following:

`TCP`	Enables DLPI driver statistics when under the TCP/IP stack.
`HMP`	Enables Hyper Messaging Protocol (HMP) statistics gathering.
`FW`	Enables firmware statistics gathering.
`RST`	Disables/resets the DLPI driver and HMP statistics.

  Note that turning on performance statistics gathering will affect the performance of HyperFabric, by increasing CPU usage and message latency (the time it takes a message to get from one point to another). It might also decrease the throughput of data. HP recommends that you do not specify this parameter unless you are trying to troubleshoot HyperFabric problems.

  You must be logged in as `root` to use this parameter.

- `-d` specifies that you want to set the level of data displayed to *`display_level`*, which is one of the following:

`NET`	Displays fabric component statistics.
`CFG`	Displays the management daemon (`clic_mgmtd`) configuration and statistics data.
`VRID`	Displays the virtual route identifier information—the VRIDs, IP addresses, switch hopcounts, and physical routes—for each HyperFabric adapter in the local node.
`ALL`	Displays all available status data.

- `-c` displays the statistics of the adapter identified by *`adapter_ID`*.
- `-n` displays statistics for the node identified by *`nodename`*.
- `-s` specifies that you want to get the status of the switch identified by *`switch_ID`* (assigned by `clic_init` or SAM). Note that this parameter is meaningful only if you enabled switch management (through the `clic_init` command).

  To determine the *`switch_ID`*, run the `clic_stat` command without specifying any parameters. This displays each *`switch_ID`* known to the system. Look through the command's output to find the ID of the switch you are interested in.
- `-?` displays the online help for `clic_stat`.

If you do not specify any of the above parameters, the current fabric map is displayed, showing the last known status of the components.

Examples of clic_stat

Some examples of using `clic_stat` are as follows:

- **Example 1**

  If the local node is `bently7`, and you want to enable DLPI driver statistics gathering, issue this command:

  `clic_stat -p TCP`

  The generated output could look like this:

```
================================================================================
Date: Sat Aug 5 16:08:14 2000

Node: bently7.corp2.com

--------------------------------------------------------------------------------

       Performance Statistics Levels - Possible perf degradation occurring:
           DLPI performance statistics enabled

================================================================================
```

- **Example 2**

If the local node is bently7, and you want to display the management daemon (clic_mgmtd) configuration and statistics data, issue this command:

clic_stat -d CFG

The generated output could look like this:

```
================================================================================
Date: Sat Aug 5 16:08:12 2000

Node: bently7.corp2.com

--------------------------------------------------------------------------------

================================================================================
          CLIC Management Global Status/Statistics
                Current Component Versions
        CLIC Management process version:         1.0
        CLIC Management API version:             1.0
        CLIC Driver version:                     1.0
                Global Management Statistics
        Node failures:                           0
        Nodes active:                            0
        Command session failures:                0
        Command active sessions:                 1
        Command total sessions:                  6
        Management mesh session failures:        0
        Management mesh active sessions:         0

        Management Global Interval Timers
        Fabric mapping interval (ms):            60000

        Performance Statistics Levels - Possible perf degradation occurring:
                DLPI performance statistics enabled

================================================================================
```

- **Example 3**

  a. HMP in the non-local failover mode

  If the local node is bently7, and you want to display the statistics for the adapter clic1 (which is on bently7), issue this command:

  clic_stat -c clic1

The generated output could look like this:

```
================================================================================
Date: Sat Aug 5 16:08:30 2000

Node: bently7.corp2.com

--------------------------------------------------------------------------------
        Adapter ID:                     clic1
        Instance Number:                1
        Adapter Type:                   4X HF2 PCI
        Firmware File:                  /opt/clic/firmware/clic_fw_hf232c
        Major Num:                      238
        Mgmt process driver handle:     5
        Version:                        1.0
        H/W Path:                       8/4/0/0
        Primary IP address:             192.0.0.9
        Uptime time:                    0 days 21 hours 1 min 3 sec

        Adapter State Flags Set - Multiple flags may be set:
                Configured and operational
                Switch mode

                 F/W Error Statistics
        Frame buffer overflow:                  0
        Receive on disabled endpoint:           0
        Invalid endpoint ID:                    0
        Invalid endpoint protection key:        0
        Interleaved gathered receive:           0
        Interleaved multi-frame bulk messages:  0
        NQ overflow:                            0
        Send NQ overflow:                       0
        Invalid slot key for NQ credit update:  0
        DLPI QOS receive buffer shortage:       0
        Link congestion events:                 0
        Max send packet retry exceeded:         0
        Link or switch failure events:          0
        Link or switch resume events:           1
        Bad route detected:                     0
        Bad optional data length:               0
        Invalid message received:               0
        CRC error:                              0
        Invalid CRC word:                       0
        Bad frame length:                       0
        Receive buffer overflow:                0
        Null packets generated on link reset:   2
```

```
        Firmware reset notification:            0
        Data corruption notification:           0
        Unsupported QOS message received:       0
        Invalid HMP VC ID:                      0
        Invalid HMP endpoint ID:                0
        Invalid HMP endpoint protection key:    0
        HMP message order violation:            0
        Packet drops:                           0
        Transmit side congestion events:        1
        Receive side congestion events:         0
                  Other Misc Statistics
        Mapping message send failures:          0
================================================================================
```

b. HMP in the local failover mode

If the local node is `bently7`, and you want to display the statistics for the adapter `clic1` (which is on `bently7`), issue this command:

`clic_stat -c clic1`

The generated output could look like this:

```
================================================================================
Date: Sat Aug 5 16:08:30 2000

Node: bently7.corp2.com

--------------------------------------------------------------------------------
        Adapter ID:                     clic1
        Card Pair:                      clic0
        Instance Number:                1
        Adapter Type:                   4X HF2 PCI
        Firmware File:                  /opt/clic/firmware/clic_fw_hf232c
        Major Num:                      238
        Mgmt process driver handle:     5
        Version:                        1.0
        H/W Path:                       8/4/0/0
        Primary IP address:             192.0.0.9
        Uptime time:                    0 days 21 hours 1 min 3 sec

        Adapter State Flags Set - Multiple flags may be set:
                Configured and operational
                Switch mode

                 F/W Error Statistics
```

```
        Frame buffer overflow:                  0
        Receive on disabled endpoint:           0
        Invalid endpoint ID:                    0
        Invalid endpoint protection key:        0
        Interleaved gathered receive:           0
        Interleaved multi-frame bulk messages:  0
        NQ overflow:                            0
        Send NQ overflow:                       0
        Invalid slot key for NQ credit update:  0
        DLPI QOS receive buffer shortage:       0
        Link congestion events:                 0
        Max send packet retry exceeded:         0
        Link or switch failure events:          0
        Link or switch resume events:           1
        Bad route detected:                     0
        Bad optional data length:               0
        Invalid message received:               0
        CRC error:                              0
        Invalid CRC word:                       0
        Bad frame length:                       0
        Receive buffer overflow:                0
        Null packets generated on link reset:   2
        Firmware reset notification:            0
        Data corruption notification:           0
        Unsupported QOS message received:       0
        Invalid HMP VC ID:                      0
        Invalid HMP endpoint ID:                0
        Invalid HMP endpoint protection key:    0
        HMP message order violation:            0
        Packet drops:                           0
        Transmit side congestion events:        1
        Receive side congestion events:         0
                  Other Misc Statistics
        Mapping message send failures:          0
================================================================================
```

- **Example 4**

  If the local node is `bently6`, and you want to display the VRIDs, IP addresses, switch hopcounts, and physical routes for each HyperFabric adapter in `bently6`, issue this command:

  `clic_stat -d VRID`

  The generated output if the nodes are connected in a Point-to-Point configuration could look like this:

```
================================================================================
Date: Sat Aug 5 16:08:12 2000

Node: bently6.corp4.com

--------------------------------------------------------------------------------

================================================================================

 Adapter :        clic0

--------------------------------------------------------------------------------

VRID    IP Address      Switch Hopcount   Route Id    Physical route

--------------------------------------------------------------------------------

  1      192.0.0.1              0                    Point to Point Connection

  2      192.0.0.2              0                    Point to Point Connection

  12     192.0.0.12             0                    Point to Point Connection

  13     192.0.0.13             0                    Point to Point Connection

================================================================================

 Adapter :        clic2

--------------------------------------------------------------------------------

VRID    IP Address      Switch Hopcount   Route Id    Physical route

--------------------------------------------------------------------------------

  1      192.0.0.1              0                    Point to Point Connection

  2      192.0.0.2              0                    Point to Point Connection

  12     192.0.0.12             0                    Point to Point Connection

  13     192.0.0.13             0                    Point to Point Connection

================================================================================
```

The generated output if the nodes are connected through a switch could look like this:

```
================================================================================
Date: Sat Aug 5 16:08:12 2000

Node: bently6.corp4.com

--------------------------------------------------------------------------------

================================================================================

 Adapter :        clic0

--------------------------------------------------------------------------------

VRID    IP Address      Switch Hopcount   Route Id    Physical route

--------------------------------------------------------------------------------

  1      192.0.0.1             1               0           0x09

  2      192.0.0.2             1               0           0x03

  12     192.0.0.12            1               0           0x06

  13     192.0.0.13            1               0           0x00

================================================================================

 Adapter :        clic2

--------------------------------------------------------------------------------

VRID    IP Address      Switch Hopcount   Route Id    Physical route

--------------------------------------------------------------------------------

  1      192.0.0.1             1               0           0x03

  2      192.0.0.2             1               0           0x3d

  12     192.0.0.12            1               0           0x00

  13     192.0.0.13            1               0           0x3a
```

```
================================================================================
```

- **Example 5**

  If the local node is `bently7`, and you want to disable all statistics gathering on `bently7`, issue this command:

  `clic_stat -p RST`

  The generated output could look like this:

```
================================================================================
Date: Sat Aug 5 16:08:35 2000

Node: bently7.corp6.com

--------------------------------------------------------------------------------
        No performance statistics are being gathered - Reset successful

================================================================================
```

Viewing man Pages

If you want to be able to view the HyperFabric man pages, you must first add `/opt/clic/share/man` to your `MANPATH` environment variable. Then, to view a man page, type the following:

man *command_name*

For example, to view the man page for `clic_stat`, type this:

```
man clic_stat
```

Stopping HyperFabric

You can stop HyperFabric only on a local node. Stopping HyperFabric on a node makes all of the HyperFabric adapters in that node unreachable by all other nodes in the fabric. It stops the HyperFabric management process, which stops all interconnect operations on the node.

To stop HyperFabric, you use (1) the `clic_shutdown` command (described below) or (2) SAM (described in "Using SAM" on page 129).

Note that if you stop HyperFabric, the only ways you can restart HyperFabric on the node is to (1) run the `clic_start` command (see "Using the clic_start Command" on page 110) or (2) use SAM (see "Using SAM" on page 110).

Using the clic_shutdown Command

Run the `clic_shutdown` command to stop HyperFabric on the local node.

If you include `/opt/clic/bin` in your `PATH` statement, you can run the command as it is shown below. Otherwise, you must include `/opt/clic/bin` as part of the command name (that is, `/opt/clic/bin/clic_shutdown`).

You must be logged in as `root` to run this command.

The syntax is as follows:

```
clic_shutdown
```

Note that you also can issue the command `clic_shutdown -?` to display the online help for `clic_shutdown`.

Using SAM

To use SAM to stop HyperFabric on a local HP 9000 system running HP-UX 11i v2, complete the following steps:

Step 1. Start SAM.

Step 2. Select the “Networking and Communications” area.

Step 3. Select “HyperFabric.”

Step 4. Pull down the “Actions” menu and select `Stop HyperFabric`. Note that if HyperFabric is not running on the system, `Stop HyperFabric` is grayed out and you cannot select it.

When HyperFabric stops, a confirmation message displays. Also, the status “HyperFabric: Not running” is displayed above the adapter configuration area of the screen.

Step 5. Exit SAM.

6 Troubleshooting HyperFabric

This chapter contains these sections that describe troubleshooting HyperFabric:

Running Diagnostics

Before running HyperFabric diagnostics:

1. Confirm HyperFabric adapters are installed on each node. Use the following command to display a list of HyperFabric adapters installed each node:

   ```
   $ ioscan -funC clic
   ```

2. Check to see if HyperFabric software is installed. Use the following command to display the HyperFabric software version installed on each node:

   ```
   $ swlist | grep -i hyperfabric
   ```

3. Check the patch level. Use the following command to display the list of patches installed on each node:

   ```
   $ swlist
   ```

   Refer to the *HyperFabric Release Notes* to determine which patches need to be installed.

4. Make sure the HyperFabric subsytem is started. Use the following command to determine if the HyperFabric subsystem is running on the node:

   ```
   $ ps -ef | grep clic
   ```

   If the HyperFabric daemon (`clic_mgmtd`) is not running, start the HyperFabric subsystem by executing the following command:

   ```
   $ /opt/clic/bin/clic_start
   ```

5. Make sure the correct IP addresses are assigned to the HyperFabric adapters. Use the following command to display all of the network interfaces and the IP addresses assigned to them:

   ```
   $ netstat -in
   ```

   If an IP address is not assigned to a HyperFabric adapter, execute the following commands in the order listed:

   ```
   $ clic_shutdown
   ```

   ```
   $ clic_init
   ```

$ clic_start

All of these commands reside in the `/opt/clic/bin` directory.

6. Check cabling to make sure all of the HyperFabric adapters are connected to the fabric.
7. Run the following command:

   $ /opt/clic/bin/clic_stat -dALL

   If a TCP/UDP/IP application is running:

   Check the *firmware file* field to make sure the same version of firmware is downloaded on all of the HyperFabric adapters in the cluster. If this is not the case, run the following commands in the order listed below:

   $ clic_shutdown

   $ clic_init

   (The answer to the interoperability question must be consistent on all nodes in the fabric.)

   $ clic_start

   Check to make sure all IP addresses have been assigned.

   Check the subnet. Every HyperFabric adapter in the fabric must be able to communicate with every other HyperFabric adapter.

   Every HyperFabric adapter in the fabric must be connected, point-to-point or via a switch.

   Run diagnostics to make sure data can be transferred on the HyperFabric adapters. Use the following command (which is detailed in the next section of this chapter):

   $ clic_diag

   If the HyperFabric subsystem is still not usable, contact your HP support representative with the diagnostics data generated using the `clic_diag` command.

   If an HMP application is running:

Make sure 4X HyperFabric adapters are installed on the nodes. The *adapter type* field indicates the type of HyperFabric adapter that is installed. HMP will only run on 4X HF2 PCI (A6386A) adapters.

Check the *firmware file* field to make sure the firmware file name for each HyperFabric adapter ends in 32c. All of the HyperFabric adapters in the fabric must have firmware files that end in 32c for HMP to run. If this is not the case, run the following commands in the order listed below:

```
$ clic_shutdown
```

```
$ clic_init
```

(The answer to the interoperability question must be consistent on all nodes in the fabric.)

```
$ clic_start
```

All of the IP addresses in the fabric must be in the same subnet. The lower 10 digits of all of the IP addresses in the fabric must be unique. Every HyperFabric adapter in the fabric must be able to communicate with every other HyperFabric adapter in the fabric.

Every HyperFabric adapter in the fabric must be connected, point-to-point or via a switch. If one of the adapters in the fabric is not connected, HMP will not be able to run.

Run diagnostics to make sure data can be transferred on the HyperFabric adapters. Use the following command (which is detailed in the next section of this chapter):

```
$ clic_diag
```

If you are running an Oracle application using HMP and traffic does not appear to be flowing after completing all of the steps listed above, check the `ORAHOME/rdbms/log/alert*log` files. The display should show `Cluster Interconnect IPC version:Oracle using HP-HMP` logged in the alert log files. If not, recompile the Oracle application to run using HMP according to the instructions provided in the *Oracle Installation and Administration Guide*.

Diagnostics can be run on many of the HyperFabric components by using the `clic_diag` command. If the HyperFabric subsystem is still not usable, contact your HP support representative with the diagnostics data generated using the `clic_diag` command described below.

The clic_diag Command

Use the `clic_diag` command to run the following diagnostics:

- Probe a specific remote node.
- Dump and format trace data.
- Set the tracing level for the HyperFabric software and the firmware.

Two versions of the `clic_diag` command are available:

- The command line version. In this version, you specify the command and the parameters you want to run (which are described below).
- The interactive version. To use this version, specify only the command name (`clic_diag`). Then, you are prompted for the same kinds of information you specify when you use the command line version.

If you include `/opt/clic/bin` in your `PATH` statement, you can run the command as it is shown below. Otherwise, you must include `/opt/clic/bin` as part of the command name (that is, `/opt/clic/bin/clic_diag`).

You must be logged in as `root` to run this command.

The command line syntax is as follows:

```
clic_diag [-r remote_node_name] [-T trace_level] [-D TCP_IP]
          [-C TCP_IP] [-B trace_buffer_size]
          [-F trace_file_size] [-d adapter_ID filename] [-?]
```

Note that some of the lines in the above syntax are indented for readability purposes only. When you actually type the command, you do not indent anything.

The command parameters are as follows:

- `-r` specifies that you want to probe a specific remote node identified by *`remote_node_name`*. The probe is done on all operational routes to the remote node.
- `-T` specifies that you want to set the tracing level for the HyperFabric software and the firmware to *`trace_level`*, which can be one or more of the following (you can specify either the keyword or the hexadecimal value):

`reset` or `0x0`	Reset tracing to the default level for all components (that is, turn off all tracing).
`send_path` or `0x0002`	Trace the HyperFabric software send path.
`packet_data` or `0x0004`	Trace the packet data.
`packet_hdr` or `0x0008`	Trace the protocol headers (TCP/UDP/IP).
`data_struct` or `0x0010`	Trace the HyperFabric software CLIC-specific data structures and events.
`rcv_path` or `0x0020`	Trace the HyperFabric software receive path.
`control` or `0x0040`	Trace the HyperFabric software control messages.
`all` or `0xffff`	Turn on all possible trace levels. Note that this will have a severe impact on performance.

  To find out what the current tracing level is, run this parameter (`-T`) without any keyword or hexadecimal value specified.
- `-D` specifies that you want to dump the trace buffers from the kernel to a user-space file and format the data using the formatter `TCP_IP`. The HyperFabric software trace data is dumped into the file `/var/adm/clic_ip_drv.trc` and formatted.
- `-C TCP_IP` operates in a way similar to `-D TCP_IP`, except that the trace data in the kernel is dumped at regular intervals to the trace file `/var/adm/clic_ip_drv.trc0`, and if that file gets full, to `/var/adm/clic_ip_drv.trc1`. So, new trace events are appended to the trace output file. This stops when tracing is turned off.

- `-B` specifies that you want the size (in bytes) of the trace buffer to be *`trace_buffer_size`*. The buffer is dynamically allocated when tracing is enabled. Also, the buffer is circular, which means that when the end of the buffer is reached, the data wraps around to the beginning of the buffer (and overwrites any previous data). If you do not specify this parameter, a default buffer is created with a size of 64k bytes.
- `-F` specifies *`trace_file_size`* as the maximum size of the trace output file. For the `-D TCP_IP` parameter, the file is `/var/adm/clic_ip_drv.trc`. For the `-C TCP_IP` parameter, the files are `/var/adm/clic_ip_drv.trc0` (and if needed, `/var/adm/clic_ip_drv.trc1`).
- `-d` dumps the memory of the adapter identified by *`adapter_ID`* to the dump file *`filename`*. If you do not specify *`filename`*, the default file is `/var/adm/clic_fw.dump`*x*, where *x* is the adapter instance number.
- `-?` displays the online help for `clic_diag`.

If you do not specify any of the above parameters, the online help for `clic_diag` is displayed.

Example of clic_diag

An example of the `clic_diag` command is shown below.

If the local node is `bently6`, and you want to confirm that all of the adapters on `bently6` are communicating with the target adapters on `bently8`, issue this command:

```
$ clic_diag -r bently8
```

The generated output could look like this:

```
CLIC_PROBE: 256 byte packets
  Source adapter id: bently6.corp4.com:clic0
  Target adapter id: bently8.corp2.com:clic1

256 bytes:     seq_num = 0.    Packet Acknowledged.
256 bytes:     seq_num = 1.    Packet Acknowledged.
256 bytes:     seq_num = 2.    Packet Acknowledged.
256 bytes:     seq_num = 3.    Packet Acknowledged.
256 bytes:     seq_num = 4.    Packet Acknowledged.
---------  CLIC_PROBE Statistics --------
5 packets transmitted, 5 packets received, 0% packet loss.
CLIC_PROBE: 256 byte packets
```

```
  Source adapter id: bently6.corp4.com:clic1
  Target adapter id: bently8.corp2.com:clic3

256 bytes:     seq_num = 0.    Packet Acknowledged.
256 bytes:     seq_num = 1.    Packet Acknowledged.
256 bytes:     seq_num = 2.    Packet Acknowledged.
256 bytes:     seq_num = 3.    Packet Acknowledged.
256 bytes:     seq_num = 4.    Packet Acknowledged.
---------  CLIC_PROBE Statistics --------
5 packets transmitted, 5 packets received, 0% packet loss.
```

Using Support Tools Manager

Use **Support Tools Manager** (**STM**) with HyperFabric to gather information about HyperFabric components and to diagnose hardware problems.

Two tools are available in STM for HyperFabric:

- The Information Tool provides information about the HyperFabric adapter, without resetting the adapter.
- The Diagnostics Tool can be used to run tests on the HyperFabric adapter; the tool reports any failures.

You can run STM in three ways:

- In the X Windows environment.
- In command line mode.
- In menu mode.

See the *Support Media User's Manual* for details about using STM.

Useful Files

When you are troubleshooting HyperFabric-related problems, you might find it useful to look at the contents of the following files:

- `/etc/rc.config.d/clic_global_conf`

  This is the global configuration file. Check it to confirm that the configuration information is correct.

- `/var/adm/clic_log`

  This is a global log file that contains a history of significant HyperFabric events. For example, it contains a history of the fabric's startup and any errors that occurred during it.

- `/var/adm/clic_log.old`

  This is the backup copy of the log file that is created when the log file grows larger than 100 Kbytes

- `/var/adm/OLDclic_log`

  This is the log file from the previous time the `clic_start` command was executed.

- `/var/adm/syslog.log`

  This is the system log file, which contains a history of events occurring on the HP 9000 system.

- `/var/adm/clic_ip_drv.trc`

  This is one of the HyperFabric software's trace files, and it is created by the `clic_diag -D TCP_IP` command.

- `/var/adm/clic_ip_drv.trc0`

  This is one of the HyperFabric software's trace files, and it is the primary file created by the `clic_diag -C TCP_IP` command.

- `/var/adm/clic_ip_drv.trc1`

  This is one of the HyperFabric software's trace files, and it is created by the `clic_diag -C TCP_IP` command when the primary trace file (`clic_ip_drv.trc0`) becomes full.

`/var/adm/clic_fw.dumpx`

This is the default file for a memory dump of an HyperFabric adapter, created when the `clic_diag -d` command is run without specifying an output file.

- `/etc/rc.config.d/netconf`

  This file contains IP-related configuration information for all of the networking adapters installed in the HP 9000.

  IMPORTANT: `clic_init` and SAM modify this file, adding some HyperFabric-related lines that end with the characters `#clic`. These lines are used by the HyperFabric software—and are not comments—so do not remove them from the file.

- `/etc/services`

  This is the system service name database.

  IMPORTANT: These two HyperFabric-related lines must be in this file:

  — `hp-clic 3384/tcp #clic management daemon`

  — `hp-clic 3384/udp #clic switch management`

  These lines are used by the HyperFabric software—and are not comments—so do not remove them from the file.

LED Colors and Their Meanings

Listed below are the possible colors (and the corresponding meaning) of the LEDs on the HyperFabric adapters and switches.

Adapter LEDs

Table 6-1 below shows the names of the LEDs on each HyperFabric adapter. Note that the LEDs on the A4920A adapter are labeled, but the labels might be hard to see when a cable is connected to the adapter.

Table 6-1 LED Names (by Adapter)

LED Name	HyperFabric Adapter
“Connected/Traffic”	A4919A
“Link”	A4920A
	A6386A
“Link Connected/Traffic”	A4921A
“Error”	All

Some of the LEDs—“Connected/Traffic,” “Link,” and “Link Connected/Traffic”—are equivalent but are labeled differently, depending on the adapter. So, their colors and meanings are the same, regardless of the adapter.

The HyperFabric adapter LED colors and meanings are as follows:

✓ If the adapter is not operational, the “Connected/Traffic”, “Link”, or “Link Connected/Traffic” LED on the adapter is off. Some of the possible reasons for this happening are the following:

— The HP 9000 is not operational.

— HyperFabric has not been started on the HP 9000.

— An adapter is installed in a slot in the HP 9000, but the cable is attached incorrectly or no cable is attached at all.

- The adapter is bad.
- The cable is bad.
- The switch port is bad (if the adapter is connected to a switch). Note that if a switch port is bad, and (for some reason) you cannot use a different port on the switch, you must replace the switch module in the HF2 switch (whichever is applicable). However, you first should try turning the switch's power off and then back on.
- The adapter is connected to a non-operational adapter in the remote node (in a node-to-node configuration).

✓ If the connection from the adapter to the corresponding switch port (if a switch is used) or the corresponding adapter in the remote node (in a node-to-node configuration) is operational, the "Connected/Traffic", "Link" or "Link Connected/Traffic" LED on the adapter shows as solid green.

✓ If data is flowing between the adapter and the switch port (if a switch is used) or the corresponding adapter in the remote node (in a node-to-node configuration), the "Connected/Traffic", "Link" or "Link Connected/Traffic" LED shows as flashing green. Note that data does not start to flow until HyperFabric initialization has occurred (see "Using the clic_init Command" on page 79).

✓ If the adapter is in an error state that requires it to be replaced, the "Error" LED on the adapter shows as solid yellow. See "Replacing a HyperFabric Adapter" on page 156 if you need to replace an adapter.

Table 6-2 below summarizes the adapter LED information in a table format.

Table 6-2 HyperFabric Adapter LED Colors and Meanings

LED	Color	Meaning	Notes
"Connected/Traffic" (A4919A and A6092A adapters) "Link" (A4920A and A6386A adapters) "Link Connected/Traffic" (A4921A adapter)	None	Adapter is not operational.	• HP 9000 is not operational. • HyperFabric is not running on HP 9000. • Adapter is installed in a slot in HP 9000, but cable is attached incorrectly or no cable is attached at all. • Adapter is bad. • Cable is bad. • Switch port is bad (if adapter is connected to a switch). If switch port is bad and you cannot use a different port on the switch, you must replace the entire HF2 switch module. (First try powering the switch off and then back on.) • Adapter is connected to a non-operational adapter in the remote node (in a node-to-node configuration).
	Solid green	Connection from adapter to switch port or remote adapter is operational.	
	Flashing green	Data is flowing between the adapter and the switch port or remote adapter.	Data does not start to flow until HyperFabric initialization has occurred.

Table 6-2 HyperFabric Adapter LED Colors and Meanings (Continued)

LED	Color	Meaning	Notes
"Error"	None	Adapter is not in an error condition.	Adapter should be operating normally.
	Solid yellow	Adapter is in an error condition.	You must replace the adapter.

HF2 Switch LEDs

The HF2 switch LED colors and meanings are explained below.

- For each “Status” LED on the switch:
  - ✓ If the card/switch module is not operating, the LED is off.
    - — For the integrated Ethernet management LAN adapter card (in the top slot) and the integrated 8-port fiber card (in the middle slot): the card can be safely removed by qualified HP personnel only.
    - — For an A6388A HF2 8-port fiber or A6389A HF2 4-port copper switch module in the expansion slot (the bottom slot): the switch module can be safely removed by you or qualified HP personnel.
  - ✓ If a fault is occurring on the card/switch module, the LED shows as solid yellow.
  - ✓ If the card/switch module passed the self-test and is operating, the LED shows as solid green.
- For the “Power A” and “Power B” LEDs on the switch:
  - ✓ If the power to the switch is off, the LEDs are off.
  - ✓ If the power to the two redundant power buses is on, the LEDs show as solid green.

    Note that if only one of the “Power” LEDs shows as solid green, the switch is still operational.
- For the “Ethernet Port Main” and “Ethernet Port Aux” LEDs on the switch:
  - ✓ If the associated Ethernet port is not connected to the Ethernet network, the LED is off.
  - ✓ If the connection from the associated Ethernet port to the Ethernet network is operational, the LED shows as solid green.
  - ✓ If data is flowing between the associated Ethernet port and the Ethernet network, the LED shows as flashing green.

❑ For each port on the cards/switch module in the switch:

✓ If the port is not operational, the LED is off. Some of the possible reasons for this happening are the following:

— A cable is not attached correctly to the port or no cable is attached at all.

— The port is connected to a non-operational adapter in an HP 9000. (See "Adapter LEDs" on page 143 or Table 6-2 on page 145 for some tips about a non-operational adapter.)

— The cable is bad.

— The port is bad. Note that if a port is bad, and (for some reason) you cannot use a different port on that card/switch module, you can replace the card/switch module. Remember, though, that the integrated 8-port fiber card can be removed by qualified HP personnel only. The A6388A and A6389A switch modules can be removed by you or qualified HP personnel. However, you first should try turning the switch's power off and then back on.

✓ If the connection from the port to the corresponding adapter in the HP 9000 is operational, the LED on the port shows as solid green.

✓ If data is flowing between the port and the corresponding adapter, the LED shows as flashing green. Note that data does not start to flow until HyperFabric initialization has occurred (see "Using the clic_init Command" on page 79).

Note that a legend listing the possible port states—disconnected, connected, and traffic—and the corresponding LED colors is printed on the front of the switch, for your reference. See Figure 3-2 on page 59 and Figure 3-3 on page 60, for the locations of the switch LEDs and the legend.

Table 6-3 below summarizes the HF2 switch LED information in a table format.

Table 6-3 HF2 Switch LED Colors and Meanings

LED	Color	Meaning	Notes
"Status"	None	The card/switch module is not operating.	Integrated Ethernet management LAN adapter card (in the top slot) & integrated 8-port fiber card (in the middle slot): the card can be safely removed by qualified HP personnel only. A6388A HF2 8-port fiber or A6389A HF2 4-port copper switch module in the expansion slot (the bottom slot): the switch module can be safely removed by you or qualified HP personnel.
	Solid yellow	A fault is occurring on the card/switch module.	
	Solid green	The card/switch module passed the self-test and is operating.	
"Power A" and "Power B"	None	Power to switch is off.	
	Solid green	Power to switch is on.	One LED solid green: switch is still operational.

Table 6-3 HF2 Switch LED Colors and Meanings (Continued)

LED	Color	Meaning	Notes
"Ethernet Port Main" and "Ethernet Port Aux"	None	Ethernet port is disconnected.	This can happen if the Ethernet port is not correctly connected to the Ethernet network, or the integrated Ethernet management LAN adapter card is experiencing a fault. Management is disabled when the Ethernet port is not operating, however, this does not prevent the fabric from operating.
	Solid green	Connection from the Ethernet port to the Ethernet network is operational.	
	Flashing green	Data is flowing between the Ethernet port and the Ethernet network.	

Table 6-3 HF2 Switch LED Colors and Meanings (Continued)

LED	Color	Meaning	Notes
"Port *x*"	None	Port is not operational.	• A cable is not attached correctly to the port or no cable is attached at all. • The port is connected to a non-operational adapter in an HP 9000. (See "Adapter LEDs" on page 143 or Table 6-2 on page 145 for some tips about a non-operational adapter.) • The cable is bad. • The port is bad. Note that if a port is bad, and (for some reason) you cannot use a different port on that card/switch module, you can replace the card/switch module. The integrated 8-port fiber card (in the middle slot): the card can be safely removed by qualified HP personnel only. A6388A HF2 8-port fiber or A6389A HF2 4-port copper switch module in the expansion slot (the bottom slot): the switch module can be safely removed by you or qualified HP personnel. First, try powering the switch off and then back on.
	Solid green	Connection from port to adapter is operational.	
	Flashing green	Data is flowing between the port and the adapter.	Data does not start to flow until HyperFabric initialization has occurred.

Determining Whether an Adapter or a Cable is Faulty

If you suspect that either an HyperFabric adapter or a cable attached to it is faulty, follow these steps to determine which component needs replacing:

Step 1. Disconnect the cable from the HyperFabric adapter.

Step 2. Attach a loopback plug to the adapter:

- For the A6386A adapter: Be sure to use a fiber loopback plug (one [HP part number A6384-67004] is shipped with each HF2 switch).

Step 3. Determine if the adapter is faulty:

- If the adapter is functioning correctly, the “Connected/Traffic,” “Link,” or “Link Connected/Traffic” LED on the adapter shows as solid green.
- If the adapter is faulty, the “Error” LED on the adapter shows as solid yellow.

See Table 6-1 on page 143 for a list of the LED names for each HyperFabric adapter.

Step 4. If the adapter is faulty, you must replace it (see “Replacing a HyperFabric Adapter” on page 156).

If the adapter is not faulty, assume that the cable is faulty and you must replace it.

Determining Whether a Switch is Faulty

If you suspect that a HyperFabric switch is faulty, follow the steps below to determine if the switch needs replacing.

HF2 Switch

The steps for determining if an HF2 switch is faulty are as follows:

Step 1. Check the cards/switch module in the switch—the integrated Ethernet management LAN adapter card, integrated 8-port fiber card, and switch module in the expansion slot:

- **a.** If the card/switch module is functioning correctly, its "Operating/Fault" LED shows as solid green.
- **b.** If the card/switch module is experiencing a fault, its "Operating/Fault" LED shows as solid yellow.
- **c.** If the card/switch module is not operating, its "Operating/Fault" LED is off.

Step 2. Check the switch's power:

- **a.** If the power to the switch is on, the "Power A" and "Power B" LEDs both show as solid green.
- **b.** If only one of the "Power" LEDs shows as solid green, the switch is still operational. However, the integrated Ethernet management LAN adapter card should be replaced soon. (Remember, it can be replaced by qualified HP personnel only.)
- **c.** If both "Power" LEDs are off, the enclosure (the power supply, backplane, and fan-monitoring board) must be replaced by qualified HP personnel only.

Step 3. Check the switch's Ethernet ports:

- **a.** If the port is connected to an operational Ethernet network, the port's LED ("Ethernet Port Main" or "Ethernet Port Aux") shows as solid or blinking green.
- **b.** If the port is disconnected, or connected to an Ethernet network that is not operational, the port's LED is off.

Step 4. Check the ports on the cards/switch module:

a. Disconnect the cable from the card/switch module port you suspect is faulty.

b. Attach a loopback plug to the relevant port:

- If the port is on the integrated 8-port fiber card or the A6388A HF2 8-port fiber switch module in the expansion slot, use a fiber loopback plug. (A fiber loopback plug [HP part number A6384-67004] is shipped with each HF2 switch).

c. Determine if the port is faulty:

- If the port is functioning correctly, the "Port *x*" LED shows as solid or blinking green.
- If the port is faulty, the "Port *x*" LED is off.

Repeat steps a through c for each port on the cards/switch module, to determine if any of them are faulty.

d. If the port is faulty, do not use it. Instead, use a known good port on that card/switch module if one is available. If no good ports are available, replace that card/switch module. Remember, though, that the integrated 8-port fiber card can be removed by qualified HP personnel only. The A6388A and A6389A switch modules can be removed by you or a qualified HP personnel.

Step 5. If a fan is not running, the fan tray must be replaced by qualified HP personnel only.

Replacing a HyperFabric Adapter

If OLAR is supported for the HP 9000 system and the HyperFabric adapter, and you want to do OLR, see “Online Replacement (OLR)” on page 49. The *HP HyperFabric Release Note*s contains information about which HP 9000 systems and HyperFabric adapters OLAR is supported for.

If you cannot or do not want to do OLR to replace a HyperFabric adapter, follow these steps:

Step 1. Stop HyperFabric on the node where the faulty adapter is installed (see “Stopping HyperFabric” on page 128).

Step 2. Depending on the type of HP 9000 the adapter is in, follow the steps in the HP 9000’s documentation for installing the adapter.

Remember that, when you replace a HyperFabric adapter, you do not have to re-install the HyperFabric software.

Step 3. Start HyperFabric on the HP 9000 (see “Starting HyperFabric” on page 109).

Replacing a HyperFabric Switch

If you need to replace a faulty HyperFabric switch, follow these steps:

Step 1. If you do not have an HA environment, stop HyperFabric on all nodes that are connected to the faulty switch. (See “Stopping HyperFabric” on page 128.)

If you have an HA environment with two switches (where the backup switch will have taken over when the first switch failed), disconnect all cables attached to the faulty switch.

Step 2. Install a new switch. (See “Installing HyperFabric Switches” on page 57.)

Step 3. If you stopped HyperFabric on all nodes connected to the switch (in step 1), start it on those nodes. (See “Starting HyperFabric” on page 109.)

A Safety and Regulatory Information

This appendix contains the following sections that contain safety information and regulatory statements for the HyperFabric hardware

components:

- “Safety Symbols” on page 161.
- “Regulatory Statements” on page 162.

These components were tested for conformance to various national and international regulations and standards. The scope of this regulatory testing includes electrical and mechanical safety, electromagnetic emissions, immunity, acoustics, and hazardous materials.

When required, approvals are obtained from third party test agencies. Approval marks appear on the product label. In addition, various regulatory bodies require some information under the headings noted below.

Safety Symbols

The safety-related symbols used in this manual are shown below.

WARNING **A WARNING denotes a hazard that can cause personal injury.**

CAUTION A CAUTION denotes a hazard that can damage equipment.

Regulatory Statements

This section contains the regulatory statements for the HyperFabric products.

Adapters and Switches

The following statements apply to the HyperFabric adapters and switches:

- FCC Part 15 Class A
- ICES-003 Class A
- CISPR 22 Class A
- EN55022 Class A
- EMC Directive 89/336/EEC
- AS/NZS 3548 Class A
- VCCI Class A (applies to switch only)
- UL and cUL Listed (applies to switch only)
- TÜV Certificated (applies to switch only)
- EN 60950
- Low Voltage Directive 73/23/EEC
- CB Certificated, IEC 950 (applies to switch only)

FCC Statement (USA only)

This device complies with part 15 of the FCC rules. Operation is subject to the following two conditions: (1) This device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

This equipment has been tested and found to comply with the limits of a Class A digital device, pursuant to part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy, and

if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Any changes or modifications not expressly approved by Hewlett-Packard could void the user's authority to operate this equipment.

Use of shielded interface cables is required to comply within the Class A limits in part 15 of the FCC rules.

DOC Statement (Canada only)

This Class A digital apparatus meets all the requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classe A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

Europe RFI Statement

This is a Class A product. In a domestic environment, this product may cause radio interference, in which case the user may be required to take adequate measures.

Australia and New Zealand EMI Statement

This product meets the applicable requirements of the Australia and New Zealand EMC Framework.

Radio Frequency Interference (Japan Only)

VCCI, Class A (Model A4891-62001 only)

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Declarations of Conformity

DECLARATION OF CONFORMITY
according to ISO/IEC Guide 22 and EN 45014

Manufacturer's Name: Hewlett-Packard Company

Manufacturer's Address: 8000 Foothills Blvd.
Roseville, CA 95747
USA

declares, that the product

Product Name: HyperFabric high-speed network link

Model Number(s): A4891-62001 Network Switch (Prod. No. A4891A)
A4919-60001 (PCI) Adapter Card (Prod. No. A4919A)
A4919-60002 (HSC) Adapter Card (Prod. No. A4920A)
A4919-60003 (HSC) Adapter Card (Prod. No. A4921A)

Product Options: All

conforms to the following Product Specifications:

Safety: IEC 950:1991 + A1, A2, A3, A4 / EN 60950:1992 + A1, A2, A3, A4

EMC: CISPR 22:1993 / EN 55022:1994 - Class A[1]
EN 50082-1:1992, Generic Immunity, including:
IEC 801-2:1991 / prEN 55024-2:1992, 4 kV CD, 8 kV AD
IEC 801-3:1984 / prEN 55024-3:1991, 3 V/m
IEC 801-4:1988 / prEN 55024-4:1993, 1 kV Power Lines
0.5 kV Signal Lines
IEC 1000-3-2:1995 / EN 61000-3-2:1995 (no limit)
IEC 1000-3-3:1994 / EN 61000-3-3:1995 (not applicable)

Supplementary Information:

The product herewith complies with the requirements of the Low Voltage Directive 73/23/EEC and the EMC Directive 89/336/EEC and carries the CE marking accordingly.

1) The Product was tested in a typical configuration with Hewlett-Packard information technology equipment.

Roseville, CA, June 12, 1998

Frank D. Dembski Jr., Quality Manager

European Contact: Your local Hewlett-Packard Sales and Service Office or Hewlett-Packard GmbH, Department TRE, Herrenberger Straße 130, D-71034 Böblingen (FAX: + 49-7031-14-3143)

DECLARATION OF CONFORMITY
According to ISO/IEC Guide 22 and EN 45014

Manufacturer's Name: Hewlett-Packard Company
Systems Interconnect Solutions Lab

Manufacturer's Address: 8000 Foothills Blvd.
Roseville, CA 95747
USA

declares, that the product

Product Name: Hyperfabric PCI 4X Adapter

Model Number(s): A6092-60001 (Product No. A6092A)

Product Options: All

conforms to the following Product Specifications:

Safety: IEC 950:1991 + A1, A2, A3, A4 / EN 60950:1992 + A1, A2, A3, A4, A11
GB 4943-1995

EMC: CISPR 22:1993 / EN 55022:1994 & A2 1996 - Class A[1]
CNS 13438, GB 9254-1988, CFR47, Part 15 Class A,
CISPR 24:1997 / EN 55024:1998
IEC 61000-4-2
IEC 61000-4-3 / ENV 50204
IEC 61000-4-4
IEC 61000-4-8

Supplementary Information:

The product herewith complies with the requirements of the EMC Directive 89/336/EEC and carries the CE marking accordingly.

1) The Product was tested in a typical configuration with Hewlett-Packard information technology equipment.

Roseville, CA, March 27, 2000

Frank D. Dembski Jr., Quality Manager

European Contact: Your local Hewlett-Packard Sales and Service Office or Hewlett-Packard GmbH, Department HQ-TRE, Herrenberger Straße 130, D-71034 Böblingen (FAX: + 49-7031-14-3143)

DECLARATION OF CONFORMITY
According to ISO/IEC Guide 22 and EN 45014

Manufacturer's Name: Hewlett-Packard Company
Systems Interconnect Solutions Lab

Manufacturer's Address: 8000 Foothills Blvd.
Roseville, CA 95747
USA

declares, that the product

Product Name: HyperFabric2 Switch

Model Number(s): A6384-62001 (Prod. No. A6384A) HyperFabric2 Switch
A6388-60001 (Prod. No. A6388A) 8 ports adapter card
A6389-60001 (Prod. No. A6389A) 4 ports adapter card
A6386-60001 (Prod. No. A6386A) PCI adapter card

Product Options: All

conforms to the following Product Specifications:

Safety: IEC 950:1991 + A1, A2, A3, A4 / EN 60950:1992 + A1, A2, A3, A4, A11
GB 4943-1995. IEC 825-1:1993/ EN60825-1:1994+A1, Class 1

EMC: CISPR 22:1997 / EN 55022:1998 - Class A
CNS 13438, GB 9254-1988, CFR47, Part 15 Class A
CISPR 24:1997 / EN 55024:1998
IEC 61000-4-2, IEC 61000-4-3 / ENV 50204, IEC 61000-4-4
IEC 61000-4-6, IEC 61000-4-8, EN 61000-3-2, EN 61000-3-3

Supplementary Information:

The product herewith complies with the requirements of the Low Voltage Directive 73/23/EEC and the EMC Directive 89/336/EEC and carries the CE marking accordingly.

1) The Product was tested in a typical configuration with Hewlett-Packard information technology equipment.

Cupertino, CA, April, 2001

Craig Stein, Hardware Engineering Manager

European Contact: Your local Hewlett-Packard Sales and Service Office or Hewlett-Packard GmbH, Department HQ-TRE, Herrenberger Straße 130, D-71034 Böblingen (FAX: + 49-7031-14-3143)

B Technical Specifications

This appendix contains the following sections that contain the technical specifications for the HyperFabric hardware components:

- “Environmental” on page 173.

Physical Attributes

The physical attributes of the HyperFabric adapters and switches are as follows:

- HF2 adapter (A6386A PCI [4X])
  - 4.2 inch (10.7 cm) height
  - 1 inch (2.5 cm) width
  - 7.1 inch (18 cm) length
  - 4.25 ounces (120 g) weight

- HF2 switch chassis (A6384A)
  - 3.47 inch (8.8 cm) height
  - 19 inch (48.2 cm) width (at flanges)

    17.18 inch (43.6 cm) width (enclosure)
  - 17.5 inch (44.5 cm) length (depth from back of flanges to back plate)

    18.5 inch (47 cm) length (depth including protrusion of the switch module extractors)
  - 200 ounces (5680 g) weight
- HF2 switch modules
  - A6388A
    - 0.8 inch (2 cm) height
    - 15.75 inch (40 cm) width
    - 11.5 inch (29.2 cm) length
    - 20 ounces (570 g) weight
  - A6389A
    - 0.8 inch (2 cm) height
    - 15.75 inch (40 cm) width
    - 11.5 inch (29.2 cm) length
    - 20 ounces (570 g) weight

Environmental

These environmental specifications are the same for the HyperFabric adapters and switches:

- Temperature
  - — -40 to +70 degrees C, non-operating/storage
  - — +5 to +40 degrees C, operating
  - — +20 to +30 degrees C, recommended operating range
- Relative Humidity
  - — 15% to 90%, non-operating/storage
  - — 15% to 80% @ 22 degrees C, operating
  - — 15% to 80% @ 22 degrees C, recommended operating range
- Altitude
  - — 15,000 feet (4.6 km), non-operating
  - — 10,000 feet (3.1 km), operating

Index

Index